ETSI TR 102 603 V1.1.1 (2009-01)

Technical Report

Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Connection Control Protocol (C2P) for DVB-RCS; Background Information

Reference
DTR/SES-00295

Keywords
broadband, DVB, multimedia, satellite

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://webapp.etsi.org/IPR/home.asp).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Technical Report (TR) has been produced by ETSI Technical Committee Satellite Earth Stations and Systems (SES).

1 Scope

The present BSM document provides background information for the specification of a Connection Control Protocol (C2P) for DVB-RCS networks [i.11], relying on specific DVB-RCS signalling.

C2P (and any other connection control protocol) is part of the control plane functionality and is generally used to establish connections with adequate resources in order to satisfy the service requirements of various multimedia applications. In this context, it can be seen as one component of the QoS end-to-end architecture, e.g. as defined in [i.3]. Recommendations for QoS specifications for DVB-RCS networks have been proposed by the SatLabs Group [i.4]; they will be considered in the definition of C2P. Moreover, SatLabs Group has also provided recommendations for the Harmonized Management and Control Plane Specifications in DVB-RCS networks [i.7]; some of the specified data structures (with enhancements) will be used in the definition of C2P specification [i.11].

To facilitate the definition of C2P, the present document provides and overview of the following aspects:

- Reference scenarios for DVB-RCS networks.
- C2P core elements and their inter-relationships, captured in various C2P models.
- Internetworking issues relevant to C2P, including the QoS models adopted for C2P and various terminal data structures (tables) for dynamic connectivity support.

The intension is to identify those elements common to all network reference scenarios, in order to define a basic framework for the specification of a connection protocol applicable to all scenarios.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific.

- For a specific reference, subsequent revisions do not apply.
- Non-specific reference may be made only to a complete document or a part thereof and only in the following cases:
  - if it is accepted that it will be possible to use all future changes of the referenced document for the purposes of the referring document;
  - for informative references.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

For online referenced documents, information sufficient to identify and locate the source shall be provided. Preferably, the primary source of the referenced document should be cited, in order to ensure traceability. Furthermore, the reference should, as far as possible, remain valid for the expected life of the document. The reference shall include the method of access to the referenced document and the full network address, with the same punctuation and use of upper case and lower case letters.

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are indispensable for the application of the present document. For dated references, only the edition cited applies. For non-specific references, the latest edition of the referenced document (including any amendments) applies.

Not applicable.

2.2 Informative references

The following referenced documents are not essential to the use of the present document but they assist the user with regard to a particular subject area**.** For non-specific references, the latest version of the referenced document (including any amendments) applies.

[i.1] ETSI EN 301 790 (V1.4.1): "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems".

[i.2] DVB-RCS 305 (06/03/01): Final Report of the AHG-RSAT Group.

[i.3] ETSI TS 102 462 (V1.1.1): "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); QoS Functional Architecture".

[i.4] SatLabs System Recommendations - Quality of Service Specifications, v1.

NOTE: Available at satlabs.org.

[i.5] ETSI EN 300 421: "Digital Video Broadcasting (DVB): Framing structure, channel coding and modulation for 11/12 GHz satellite services".

[i.6] ETSI EN 302 307 (V1.1.1): "Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications".

[i.7] SatLabs System Recommendations - Management & Control Planes Specifications, v1.

NOTE: Available at satlabs.org.

[i.8] ETSI TS 102 429-1 (V1.1.1): "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM); Regenerative Satellite Mesh - B (RSM-B); DVB-S/DVB-RCS family for regenerative satellites; Part 1: System overview".

[i.9] ETSI TS 102 293 (V1.1.1): "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP Interworking over satellite; Multicast group management; IGMP adaptation".

[i.10] ETSI TS 102 294 (V1.1.1): "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia (BSM) services and architectures; IP interworking via satellite; Multicast functional architecture".

[i.11] ETSI TS 102 602 (V1.1.1): "Satellite Earth Stations and Systems (SES); Broadband Satellite Multimedia; Connection Control Protocol (C2P) for DVB-RCS; Specifications".

[i.12] ETSI EN 301 192 (V1.4.2): "Digital Video Broadcasting (DVB); DVB specification for data broadcasting".

[i.13] ETSI TS 102 606 (V1.1.1): "Digital Video Broadcasting (DVB); Generic Stream Encapsulation (GSE) Protocol".

[i.14] IETF RFC 2474: "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers".

[i.15] IETF RFC 2475: " An Architecture for Differentiated Service".

[i.16] IETF RFC 1633: "Integrated Services in the Internet Architecture: an Overview".

[i.17] IETF RFC 2212: "Specification of Guaranteed Quality of Service".

[i.18] IETF RFC 2215: "General Characterization Parameters for Integrated Service Network Elements".

[i.19] IETF RFC 2205: "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification".

[i.20] IETF RFC 4609: "Protocol Independent Multicast - Sparse Mode (PIM-SM) Multicast Routing Security Issues and Enhancements".

[i.21] IETF RFC 826: "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addressed to 48.bit Ethernet Address for Transmission on Ethernet Hardware".

[i.22] IETF RFC 2237: "Japanese Character Encoding for Internet Messages".

[i.23] IETF RFC 4080: "Next Steps in Signalling (NSIS): Framework".

[i.24] IETF RFC 2362: "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification".

[i.25] IETF RFC 2236: "Internet Group Management Protocol version 2 (IGMPv2)".

[i.26] IETF RFC 3376: "Internet Group Management Protocol version 3 (IGMPv3)".

[i.27] AmerHis System, Interactive Broadband DVB-RCS/S OBP Communication System.

NOTE: Available at telecom.esa.int.

[i.28] SkyplexNet System.

NOTE: Available at telecom.esa.int.

[i.29] SpaceMux.

NOTE: Available at space.gc.ca or telesat.ca.

[i.30] B. Le Stradic, M. Vaissiere, 0. Bowles, Matra Marconi Space, United Kingdom "The WEST Project: exploiting the Ka-band spectrum to develop the Global Information Infrastructure" September 15 - 18, 1997 Sorrento, Italy, 3rd Ka-Band Utilization Conference.

[i.31] IETF RFC 2774: "An HTTP Extension Framework".

[i.32] IETF RFC 2211: "Specification of the Controlled-Load Network Element Service".

[i.33] IETF RFC 2597: "Assured Forwarding PHB Group".

[i.34] IETF RFC 3246: "An Expedited Forwarding PHB (Per-Hop Behavior)".

[i.35] ETSI TR 101 790 (V1.3.1): "Digital Video Broadcasting (DVB); Interaction channel for satellite distribution systems; Guidelines for the use of EN 301 790".

[i.36] IETF RFC 1518: " An Architecture for IP Address Allocation with CIDR".

[i.37] IETF RFC 1519: "Classless Inter-Domain Routing (CIDR); an Address Assignment and Aggregation Strategy".

[i.38] ITU-T Recommendation Q.2931: "Digital Subscriber Signalling System No. 2 - User-Network Interface (UNI) layer 3 specification for basic call/connection control; Modified by ITU-T Q.2971 (10/1995)".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of the present document, the following terms and definitions apply.

Class of Service (CoS): specific behaviour regarding traffic handling/forwarding

NOTE: Can be used to categorize the traffic into different classes.

connection: layer 2 logical association between two or more network entities characterized by a C2P Class of Service (C2P CoS)

Connection Control Protocol (C2P): protocol that provides the interaction between RCSTs and NCC to support the set-up, modification and release of connections

control plane: the control plane in a layered protocol architecture; it performs, among others, connection control functions, including the signalling necessary to set up, supervise and release connections

Digital Video Broadcasting Return Channel by Satellite (DVB-RCS): protocol for an interaction (or return) channel in satellite links

Digital Video Broadcasting via Satellite (DVB-S): protocol for broadcasting TV signals and, by extension, data over satellite

gateway: general term to identify both the TSGW and the RSGW

interactive network: group of terminals serviced by an NCC

IP flow: sequence of IP packets from an IP source to an IP destination

NOTE: An IP flow may be identified based on the following attributes: IP source and destination address, layer 4 protocol type, source and destination ports, class of service, router or switch interface.

management plane: in a layered protocol architecture, it provides two types of functions, namely layer management and plane management functions

Management Station (MS): network element that manages all the elements of the system of one satellite interactive network (IN); it also controls the sessions, resources and connections of the ground terminals; it is composed of the NMC and the NCC

mesh connection: direct connection established between two RCSTs

multicast: communication capability, which denotes unidirectional distribution from a single source access point to a number of specified destination access points

Network Control Centre (NCC): network element that provide real time control of the IN (e.g. session control, connection control, routing, terminals' access control to satellite resources, etc.)

Network Management Centre (NMC): network element in charge of the management of all the system elements in the IN

Network Operation Centre (NOC): is responsible for the centralized management and control functions in systems supporting multiple interactive networks, each controlled by its own NCC; NOC provides service and network (bandwidth) provisioning to the interactive network, co-ordination between NCCs, etc.

NOTE: In case of single interactive network the NOC and NCC functionality are merged.

On-Board Processor (OBP): router or switch or multiplexer in the sky; it can decouple the uplink and downlink air interface formats (modulation, coding, framing, etc.)

Quality of Service (QoS): network ability to provide service differentiation/guarantees and thus influence the perceived quality of communications with regard to a number of parameters (including delay, jitter, packet loss) that packets sent by the application experience when being transferred by the network

Return Channel Satellite Terminal (RCST): network element that provides the interface between the satellite system and external users

Regenerative Satellite Gateway (RSGW): network element in a regenerative satellite system that provides interconnection with terrestrial networks (Internet, ISDN/POTS and Intranet)

star connections: connections involving a gateway (TSGW in a transparent system or RSGW in a regenerative system).

NOTE: Star connections can involve one hop or double hop.

stream: logical flow of layer 2 data from one network reference point into the satellite network, resulting from the encapsulation of IP datagrams into MAC packets

Transparent Satellite Gateway (GW/TSGW): network element in a transparent satellite system that provides interfaces with terrestrial networks (Internet, ISDN/POTS and Intranet)

NOTE: The GW is typically integrated with the NCC in a single network element denoted as Hub.

user plane: user plane in a layered protocol architecture that provides the transfer of user data, along with associated controls (e.g. flow control, recovery from errors, etc.)

3.2 Abbreviations

For the purposes of the present document, the following abbreviations apply:

ActiveCnx Entry	Active Connection (Table) Entry
ActiveCnx Index	Active Connection (Table) Index
ADC	Analogue-to-Digital Conversion
ADR	Average Data Rate
AF	Assured Forwarding (DiffServ PHB)
ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
AVBDC	Absolute Volume Based Dynamic Capacity
AVBDCRepTime	AVBDC Repetition Time
BE	Best Effort (DiffServ PHB)
BoD	Bandwidth on Demand
BSM	Broadband Satellite Multimedia
C2P CoS	C2P Class of Service
C2P PDR	C2P Peak Data Rate (used in C2P request messages)
C2P SDR	C2P Sustainable Data Rate (used in C2P request messages)
C2P	Connection Control Protocol
Channel_ID	Channel Identifier
Channel_ID_NCC	Channel Identifier at NCC
Channel_IDxy	Channel identifier for MAC CoS y in the connectivity channel x
CL	Controlled Load (IntServ Class of Service)
CMT	Correction Map Table
Cnx PDR'	(Admitted) Connection Peak Data Rate (used in C2P response messages)
Cnx SDR'	(Admitted) Connection Sustainable Data Rate (used in C2P response messages)
Cnx	Connection
CnxProfile Entry	Connection Profile (Mapping Table) Entry
CnxProfile Index	Connection Profile (Mapping Table) Index
CnxRef ID	Connection Reference Identifier (used in the Active Connection Table)
CoS	Class of Service
CPU	Computer Processing Unit
CR	Capacity Request
CRA	Constant Rate Assignment
CRC	Cyclic Redundancy Check
CSC	Common Signalling Channel
DAC	Digital-to-Analogue Conversion
DAMA	Demand Assignment Multiple Access
DiffServ	Differentiated Services
DL	Down Link
DR	Designated Router
DRM	Digital Rights Management (for security)
DSCP	Differentiated Service Code Point
DSM-CC	Digital Storage Medium - Command and Control
DULM	Data Unit Label Method
DVB	Digital Video Broadcasting
DVB-RCS	Digital Video Broadcasting Return Channel Satellite
DVB-S	Digital Video Broadcasting by Satellite
DVB-S2	Digital Video Broadcasting by Satellite Transmission 2nd Generation
EF	Expedited Forwarding (DiffServ PHB)
ETSI	European Telecommunications Standards Institute
FCA	Free Capacity Assignment

FL	Forward Link
Group_ID	Group IDentifier
GS	Guaranteed Service (IntServ Class of Service)
GSE	Generic Stream Encapsulation
GW	GateWay
ID	IDentifier
IE	Information Element
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IMUX	Input MUltipleXer
IN	Interactive Network
INM	Inter-Network Management
IntServ	Integrated Services
IP CoS	IP class of service
IP PDR	IP Peak Data Rate (of an IP flow/flow aggregate)
IP SDR	IP Sustainable Data Rate (of an IP flow/flow aggregate)
IP	Internet Protocol
IP@	IP address
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
ITU	International Telecommunications Union
Kbps	Kilo bits per second (thousands of bits per second)
LAN	Local Area Network
LNM	Local Network Management
Logon_ID	Logon IDentifier
M&C	Management and Control
MAC CoS	MAC Class of Service
MAC	Medium Access Control
MAC@	MAC address
MAC@dst	MAC address of the destination
MAC@src	MAC address of the source
Mbps	Mega bits per second (millions of bits per second)
MCD	Multi-Carrier Demodulator
MCU	Multi-Conference Unit
MDR	Minimum Data Rate
MF-TDMA	Multiple-Frequency Time-Division Multiple Access
MIB	Management Information Base
MPE	Multi-Protocol Encapsulation
MPEG	Moving Picture Experts Group
MS	Management Station
NCC	Network Control Centre
NMC	Network Management Centre
NOC	Network Operation Centre
NSIS	Next Step In Signalling (protocol)
OAM	Operation, Administration and Maintenance
OBP	On Board Processor
OMUX	Output Multiplexer
OSI	Open System Interconnection
PDR	Peak Data Rate
Phb Entry	PHB (Mapping Table) Entry
Phb Index	PHB (Mapping Table) Index
Phb PDR'	PHB Peak Data Rate (admitted rate; used for the configuration of IP mechanisms)
Phb SDR'	PHB Sustainable Data Rate (admitted rate; used for the configuration of IP mechanisms)
PHB	Per Hop Behaviour
PID	Program IDentifier
PIM-SIM	Protocol Independent Multicast-Sparse Mode (routing protocol)
PktClass Entry	Packet Classification (Table) Entry
PktClass Index	Packet Classification (Table) Index
PSTN	Public Switched Telephone Network
PTM	Point-To-Multipoint

PTP	Point-To-Point
PVC	Permanent Virtual Circuit
QID	Queue IDentifier
QoS	Quality of Service
R1	First Rate (of a token bucket)
R2	Second rate (of a token bucket)
RBDC	Rate Based Dynamic Capacity
RBDCMax	RBDC Maximum (parameter)
RBDCTimeout	RBDC Timeout (parameter)
RC Entry	Request Class (Table) Entry
RC Index	Request Class (Table) Index
RC	Request Class
RCST	Return Channel Satellite Terminal
RF	Radio Frequency
RFC	(IETF) Request For Comments
RL	Return Link
Route_ID	Route IDentifier
RSAT	Regenerative SATellite
RSGW	Regenerative Satellite GateWay
RSM-B	Regenerative Satellite Multimedia system Family B
RSPEC	Request SPECification
RSVP	Resource reSerVation Protocol
Rx	Receive
SatLabs	Satellite Laboratories
SCD	Single Carrier Demodulator
SCPC	Single Channel Per Carrier
SDP	Session Description Protocol
SDR	Sustainable Data Rate
SIAF	Satellite Independent Access Function
S-IGMP	Satellite Internet Group Management Protocol
SIP	Session Initiation Protocol
SI-SAP	Satellite Independent Service Access Point
SLA	Service Level Agreements
SNMP	Simple Network Management Protocol
SP	Service Provider
SSR	SatLabs Systems Recommendations
ST	Satellite Terminal (e.g. in BSM)
Stream_ID	Stream IDentifier
SYNC	SYNChronization (burst type)
TBTP	Terminal Burst Time Plan
TCA	Traffic Conditioning Agreement
TCP	Transmission Control Protocol
TDM	Time Division Multiplex
TIM	Terminal Information Message
TIMb	TIM broadcast
TIMu	TIM unicast
TOS	Type Of Service
TRF	TRaFic (burst)
TS	Transport Stream
TSGW	Transparent Satellite GateWay
TSPEC	Traffic SPECification
Tx	Transmit
UDP	User Datagram Protocol
uimsbf	unsigned integer most significant bit first
UL	UpLink
VBDC	Volume Based Dynamic Capacity
VBDCMax	VBDC Maximum (parameter)
VBDCMaxBacklog	VBDC Maximum Backlog (parameter)
VBDCTimeout	VBDC Timeout (parameter)
VCC	Virtual Channel Connection
VCI	Virtual Connection Identifier
VLAN	Virtual LAN

VoIP	Voice over IP
VPI	Virtual Path Identifier
VSN	Virtual Satellite Network

4 DVB-RCS Network Reference Scenarios

Network reference scenarios are defined on the basis of a well-defined set of driver factors concerning satellite network architecture, satellite network topology and satellite payload features. In particular, the reference scenarios considered in the present document are classified with regard to the following criteria:

- Number of spot beams (network architecture):
  - single-beam;
  - multi-beam.
- Network topology:
  - star;
  - mesh.
- Satellite payload architecture:
  - transparent (or transponded);
  - regenerative.

Based on the above criteria, the following network reference scenarios are identified:

- single-beam/multi-beam, star transparent;
- single-beam/multi-beam, mesh transparent;
- single-beam/multi-beam, star/mesh regenerative.

In the above list the term "transparent" refers both to a pure transparent satellite transponder (bent-pipe) and also to a satellite transponder which does not perform any demodulation, but is able to perform some form of physical layer switching, e.g.:

- by switching uplink carriers in a beam to different downlink beams, by means of an Input Multiplexer (IMUX) and an Output Multiplexer (OMUX); or
- by performing Analogue-to-Digital Conversion (ADC) on the uplink carriers, followed by digital switching and Digital-to-Analogue Conversion (DAC).

The network reference scenarios described in this clause apply primarily to the DVB-RCS networks - a particular implementation of the Broadband Multimedia Satellite (BSM) networks. Many architectural/topological aspects of the reference scenarios also apply to the BSM networks in general; however, the description of scenarios in the present document is in terms of DVB-RCS specifics (especially with regard to the satellite air interface/waveforms).

4.1 Transparent Satellite Network

Two different transparent scenarios can be considered based on the network topology: star or mesh. They will be reviewed in the next sub-clauses. The main features of each topology will be described for both single-beam and multi-beam network architectures and the differences between them will be identified.

4.1.1 Star Transparent

In a star transparent network the communication between Return Channel Satellite Terminals (RCSTs) and Network Control Centre (NCC)/Gateway (GW) is based on a transparent (transponded) satellite. The Gateway operating in a transparent satellite network is sometimes referred to as Transparent Satellite Gateway (TSGW). The overall architecture for the star transparent network, as applicable to an Interactive Network (IN), is depicted in figure 4.1.

The network configuration includes the following elements.

- The Network Control Centre (NCC), which controls the DVB-RCS Interactive Network. It provides the unique reference for the system clock and generates control and timing signals for the operation of the satellite Interactive Network, to be transmitted by one or several Feeder Stations. The NCC also serves the satellite access requests from RCSTs, via management and control plane functions.

- The Feeder, that transmits a standard digital video broadcast signal (DVB-S [i.5] or DVB-S2 [i.6]) on the uplink forward link. The signal carries control and timing signalling needed for the operation of the satellite Interactive Network, possibly multiplexed with user data. The Feeder is customarily a functional building block of the NCC, therefore within the current document the Feeder (Feeders in the case of multi-beam architectures) will be considered part of the NCC.

- Traffic Gateway (GW), that provides user plane transport functions for the satellite network, i.e. it transmits user data and signalling on the forward link in DVB-S/S2 formats and receives return channel user data and signalling in DVB-RCS format. The GW also provides interfaces with the external networks, thus enabling the internetworking (via the backbone) primarily with Internet and Internet Service Provider (SP) networks, but also with the telephony oriented networks such as PSTN or ISDN.

DVB-RCS defines an Interactive Network as a group of RCSTs, an NCC, several GWs and a Feeder. There will be one unique NCC per Interactive Network, but there could be several GWs, all controlled by the same NCC; each GW will provide external access to different RCST populations (groups) within the Interactive Network. Typically there is just one GW per Interactive Network, co-located with the NCC, and this is the case considered in the present document. If several GWs are present in the satellite system, they are customarily associated with different interactive networks. In order to simplify the description and better concentrate on the network control aspects, the NCC and the GW will be considered in the following clauses of the document as one network element (NCC/GW, also referred to as Hub), with two collocated functional building blocks - the NCC and the GW. The NCC is in charge of executing control and management plane functions for the interactive network, while the GW is in charge of executing user plane functions.

As a further refinement, the NCC is often more precisely defined as one part of the Management Station (MS) which can be logically divided into two modules:

- the Network Control Centre, that provides the real time control of the interactive network, as described above (i.e. control plane functions); and the

- the Network Management Centre (NMC), in charge of element and network management functions (management plane functions).

It is expected that the NCC/GW is able to provide service differentiation/guarantees to subscribers based on a defined Quality of Service (QoS) model (e.g. DiffServ-based). This implies user plane functions (i.e. traffic conditioning/forwarding) and M&C plane functions (e.g. the configuration/control of the DiffServ mechanisms and other QoS-aware components) [i.3] and [i.4].

- The DVB-RCS Terminals (RCSTs), which provide the interface with the end users equipment. The RCSTs transmit bursts according to the DVB-RCS air interface standard [i.1] and receive a forward link based on the DVB-S/S2 air interface standard [i.5] and [i.6]. It is expected that the RCSTs are also able to provide service differentiation/guarantees based on the same QoS model (e.g. DiffServ) envisaged for the NCC/GW, and consequently the NCC will have to configure/control the DiffServ mechanisms (and MAC mechanisms) in RCSTs as well.

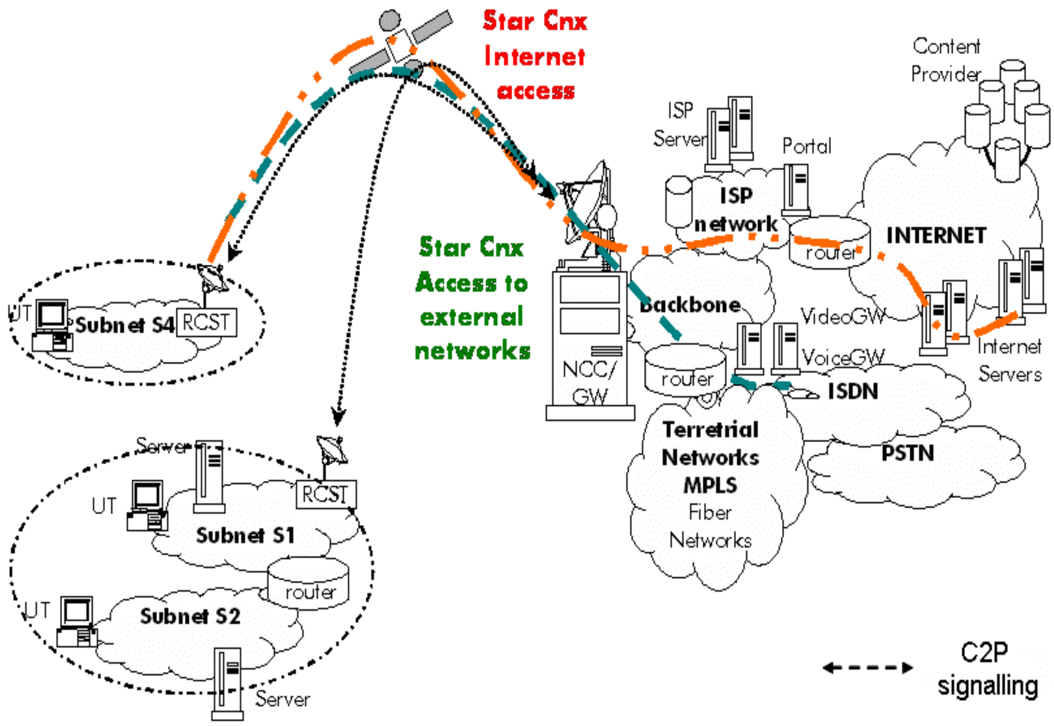

Figure 4.1: Star Transparent Network Reference Scenario

If the satellite system supports multiple Interactive Networks, each of them is under the control of its own NCC and will have its own GW. The co-ordination between NCCs is a task associated with the Network Operations Centre (NOC), which is responsible for the centralized management and control functions, such as service and network resources (bandwidth) provisioning to the interactive network providers through the NOC-NCC interface, and for the satellite configuration via the interface with the Satellite Operation Centre (SOC). In the case of a single interactive network the NOC and NCC functionality are merged.

4.1.1.1 Mobility Impact/Continuous Carrier Return Link

The DVB-RCS standard is currently being changed to provide support for mobility. For some mobile scenarios, implying many users per terminal (e.g. airplanes, ships, fast trains), the Continuous Carrier (CC) operation (also referred to as Single Channel Per Carrier (SCPC) operation) has proved to be a viable alternative to the standard MF-TDMA based return link. The last (draft) version of the DVB-RCS standard (V1.5.1) therefore includes provisions for a CC-based return link, to be supported by terminals along with the MF-TDMA return link. A reference to DVB-RCS return link could therefore imply an MF-TDMA return link, a CC return link or both. Nevertheless the MF-TDMA will be considered the default return link in the present document, but this does not preclude the use of CC for some mobile applications, where appropriate.

Support for C2P is a desirable feature not only in the conventional networks, but also in the mobile networks. DVB-RCS mobility has so far only been considered for access networks, such as those corresponding to the star transparent scenario in this clause, but it will likely be considered for other scenarios as well (e.g. mesh scenarios). While C2P is defined and specified primarily for MF-TDMA return links, the implications from using CC-based return links on various aspects of C2P (e.g. definition of C2P core elements and models, C2P QoS, etc.) will be explicitly highlighted in the present document where appropriate.

NOTE: The transmission format for CC will likely be based on DVB-S2, probably with some adaptations.

4.1.1.2 Single Beam

In the case of the single-beam scenario the RCSTs and the NCC/GW are located in the same satellite spot beam. All RCSTs belonging to an Interactive Network (IN) communicate with the NCC/GW relevant to that IN. In particular:

- each RCST transmits towards the NCC/GW by using a DVB-RCS air interface;
- the NCC/GW transmits towards the RCST by using a DVB-S/S2 air interface.

Multiple INs may be supported by the single-beam satellite system. In this case multiple NCC/GWs are envisaged, one per IN, all located in the same spot beam coverage (multi-star configuration). The whole population of RCSTs is logically divided in separated groups, one per IN. RCSTs belonging to the same group can communicate only with the NCC/GW relevant to that group. Neither single-hop communication between RCSTs belonging to the same group nor single-hop communication between RCSTs belonging to different groups is supported.

4.1.1.3 Multi Beam

A transparent multi-beam scenario implies several beams, therefore the RCSTs and possible the GWs are located in different beams. The NCC should be capable of listening and transmitting to all RCSTs in all beams, all belonging to the same IN.

A number of multi-beam architectures are possible, depending on the transparent beam connectivity supported by the satellite payload. Two examples include: 1) multiple user spot beams for the terminals which are overlapped by a single NCC/GW beam, and 2) multiple user spot beams that are transparently connected by the satellite payload to a similarly sized NCC/GW beam (i.e. a spot beam). In the latter case, the NCC/GW and the terminals are in general in different beams, except for the one terminal beam which overlaps with the NCC/GW beam. In all cases, the NCC/GW operates in the feeder link portion of the beam spectrum, distinct from the terminal beam service link spectrum.

4.1.2 Mesh Transparent

A mesh transparent satellite network is based on an architecture that supports both hub-spoke and peer-to-peer (spoke-to-spoke) communications using a transparent (transponded) satellite. The satellite payload may include a non-regenerative on board processing unit, providing channel handling and (in case of multi-beam architecture) switching among beams. The overall architecture for a star transparent system is depicted in figure 4.2.

The same system elements identified for the star transparent network are also included in the mesh transparent network. The NCC/GW performs control and management functions, similar to those described for the star transparent scenario. The NCC /GW transmits user data and signalling on the forward link using a DVB-S/S2 air interface and receives return channel user data and signalling in DVB-RCS format.

The main difference between the star transparent and mesh transparent scenarios concerns the RCSTs. In the mesh scenario the RCSTs (called mesh RCSTs) should have the capability of receiving DVB-RCS return signals (transmitted by other RCSTs), in addition to the DVB-S/S2 forward signals transmitted by the NCC/GW, in order to allow single-hop mesh communication between RCSTs. These RCSTs includes MF-TDMA burst demodulators, in addition to the DVB-S/S2 demodulators, and will be referred to as mesh terminals, to differentiate them from the standard terminals which only include DVB-S/S2 demodulators.

NOTE: Standard terminals can also be used in star transparent scenarios and star/mesh regenerative scenarios.

As for the star transparent scenario, several GWs are only included in the mesh satellite system if multiple interactive networks are supported.

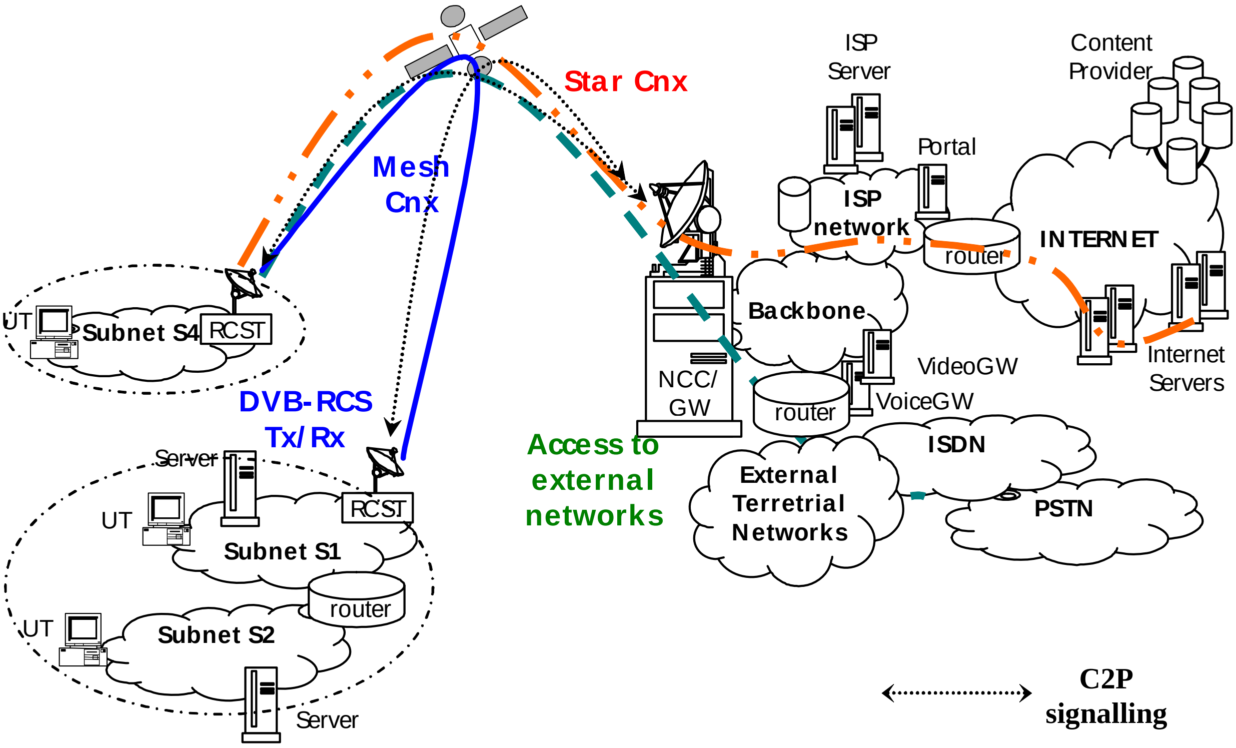

Figure 4.2: Mesh Transparent Network Reference Scenario

4.1.2.1 Single Beam

The key functional elements of this architecture are the same as in the case of star transparent networks, i.e. the RCSTs and NCC/GW, with the RCSTs having the additional capability of receiving DVB-RCS bursts on TDMA carriers. The NCC/GW functionality remains in principle the same, but additional functions will be needed for the management and control of the two receivers in the RCST (i.e. the DVB-RCS receiver and DVB-S/S2 receiver) and for the control of mesh channels and connections. The return link management information from RCSTs will be forwarded to the NCC, as in the case of the transparent star network.

In a single-beam network scenario, all RCSTs and the NCC/GW are located in the same satellite spot beam. All RCSTs belonging to an Interactive Network (IN) can in principle communicate not only with the NCC/GW relevant to that IN, but also with other mesh RCSTs. In particular:

- each RCST transmits towards the NCC/GW by using a DVB-RCS air interface;
- the NCC/GW transmits towards the RCSTs by using a DVB-S/S2 air interface;
- an RCST transmits towards other RCSTs by using a DVB-RCS air interface (the receiving RCSTs is equipped with DVB-RCS demodulators).

If the DVB-RCS demodulator is a multi-carrier demodulator (MCD), any RCST in the beam can establish mesh communications with any other RCST in the same beam. This would correspond to a fully meshed network. If the DVB-RCS demodulator is a single-carrier demodulator (SCD), an RCST can only receive transmissions from other RCSTs on the frequency on which its SCD is tuned. This may limit the connectivity between the RCSTs in the beam, leading to the concept of partially meshed network, i.e. a network that is not fully meshed. In this case the mesh network consists of a number of smaller, partially meshed subnetworks.

Multiple INs may be supported by the same single-beam, satellite system. In this case multiple NCC/GWs are envisaged, one per IN, all located in the same spot beam coverage. The whole population of RCSTs is logically divided in separated groups, one per IN. The RCSTs in a group can establish one-hop communications only with the RCSTs in the same group and with the NCC/GW for the corresponding IN.

4.1.2.2 Multi Beam

The same considerations made for the single-beam scenario also apply to the multi-beam scenario. In the latter case the RCSTs of a given IN may be located in the same or in different spot beams. Mesh communication between RCSTs in different beams assumes on board connectivity between beams. In the case of a transparent satellite the connectivity is provided at the physical layer and consists primarily in routing ("channelizing") MF-TDMA carriers or group of carriers in an uplink beam to different downlink beams.

Since the mesh transparent network is an overlay of the star network, the RCSTs are always capable of communicating with the NCC/GW of the corresponding IN. The NCC/GW should provide forward links (i.e. TDM carriers) to each downlink beam for FL signalling and also for FL access traffic, if any.

4.2 Regenerative Satellite Network

A regenerative satellite network can be configured to support both star and mesh topologies, based on DVB-RCS specification for the return channel and DVB-S/S2 specification for the forward channel. For this reason the star and mesh scenarios for the regenerative satellite payload will be jointly addressed in the next sub-clause; the differences between the two topologies will be highlighted whenever deemed necessary.

As for the transparent satellite network, the regenerative satellite network in star topology supports access traffic to/from a Gateway. The regenerative payload architecture features some enhanced network access flexibility, such as interconnecting terminals to multiple gateways and/or multiplexing star traffic with mesh traffic on the same downlink TDM carrier in a given destination beam.

4.2.1 Star/Mesh Regenerative

In a star/mesh regenerative satellite network the on-board processor capability of the satellite payload allows for both star and mesh (single-hop) communications. The overall architecture of a star/mesh regenerative system is depicted in figure 4.3. It applies to both the single-beam and multi-beam scenarios, which will be separately addressed in the next sub-clauses.

The system configuration includes the following elements:

- The On-Board Processor (OBP), which provides regeneration functions that decouple the uplink and downlink air interface formats (modulation, coding, framing, etc.). A DVB-RCS air interface is used on the uplink while a DVB-S/S2 air interface is used on the downlink. Different architectures of the OBP can be envisaged, depending on:
  - the DVB-RCS profile, which can be:
    - MPEG-based;
    - ATM-based;
  - the types of OBP functions (in addition to demodulation/modulation, decoding/coding), i.e.:
    - multiplexing (e.g. SkyPlex [i.28], SpaceMux [i.29] processors);
    - table-based switching (e.g. AmerHis [i.27] processor);
    - label-based switching (or destination-directed packet switching);
    - routing (at layer 2 or layer 3);
    - encapsulation (MPE [i.12], GSE [i.13]);
    - congestion control;
    - scheduling;
    - network time reference.

In the current implementations the routing/switching plan is usually based on connectivity/multiplexing tables, which can be statically or dynamically configured.

The OBP type may have impact on the definition and usage of various entities/elements defined in association with C2P operation (see clause 5).

The Quality of Service (QoS) support: some OBP may discriminate between different types of traffic and may assign buffer and communications resources according to the QoS requirements associated with them.

The OBP QoS support should not have any influence on the control plane functions associated with C2P, so that the same definition of the network reference scenarios may be applied to OBPs with and without QoS support.

The RCSTs, which provide interfaces with the end-users equipment and users' access to the satellite network resources. The RCSTs are able to communicate:

- with one-another in single hop, by transmitting in DVB-RCS format and receiving in DVB-S/S2 format; and
- with the Regenerative Satellite Gateway (RSGW) (see below).

The Management Station (MS), which controls the DVB-RCS Interactive Network and includes in general two main subsystems:

- the Network Control Centre (NCC), which provides control plane functions. It controls the Interactive Network and serves the satellite access requests from the RCSTs subscribers; it also manages the OBP configuration;
- the Network Management Centre (NMC), which provides management plane functions. It is in charge of element and network management functions.

For some network scenarios (e.g. RSM-B systems) the MS include a third subsystem:

- the NCC-RCST (or MS-RCST), which supports the modulation and demodulation functions, transmitting in DVB-RCS format and receiving in DVB-S/S2 format, i.e. from the point of view of the air interface/RF equipment the NCC (or MS) in the RSM-B systems behaves like a standard RCST;
- the Regenerative Satellite Gateway (RSGW), which provides RCSTs' internetworking capabilities to external networks such as Internet and PSTN or ISDN. From the RCST's point of view the RSGW provides the user plane transport functions of the Hub used in the star transparent networks (clause 4.1.1).

For some network scenarios (e.g. RSM-B systems) the RSGW include a second element:

- The RSGW-RCST, similar to the NCC-RCST. In other words, from the point of view of the air interface/RF equipment, the RSGW in the RSM-B system behaves like a large RCST.

In such scenarios the RSGW can be seen as a low cost station that may attract medium to small Service Providers, offering various services such as Internet/Intranet access, voice and videoconferencing, multicast streaming transmission/reception and PSTN/ISDN access. The RSGW is designed to share its satellite and terrestrial bandwidth resources among a relatively large number of simultaneously active subscribers (typically in the order of hundreds). In this sense the RSGW differs from a standard RCST that may support communications on behalf of a rather reduced number of subscribers (typically about ten), due to limitations regarding IP routing and the maximum number of PIDs supported for transmission and reception.

Similar to the GW used in transparent network scenarios, the RSGW is expected to provide QoS support (service differentiation and/or QoS guarantees, traffic shaping), by adopting different QoS models/policies and different subscription levels.

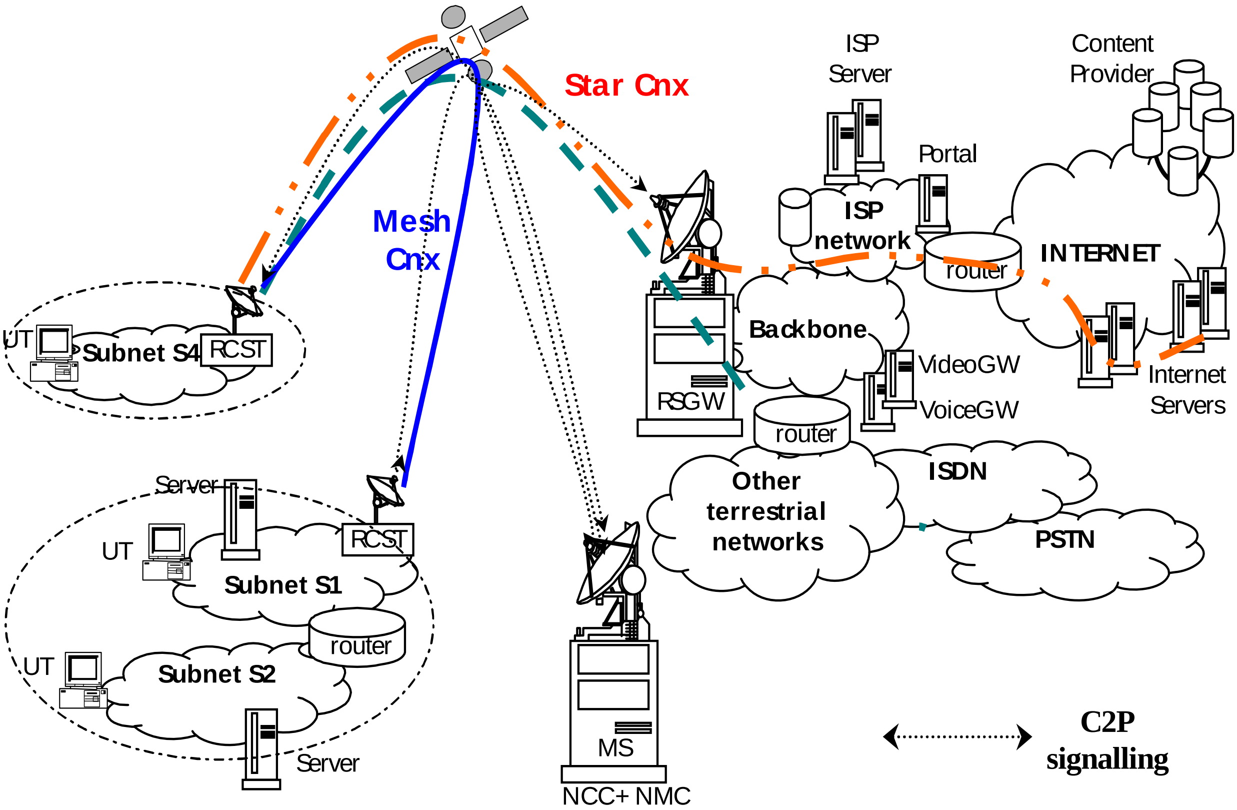

Figure 4.3: Star/Mesh Regenerative Network Reference Scenario

Future network scenarios may consider other air interfaces for the RSGW, in line with the current trends in the evolution of the DVB-RCS standard (e.g. the use of DVB-S2 on the uplink). Such air interfaces are left for future revisions of the present document, but it is expected that they will have no impact on the C2P basic requirements and only minor impact on C2P detailed specification (perhaps related to implementation options).

4.2.1.1 Single Beam

In the case of single-beam scenario all RCSTs, the RSGW(s) and the MS are located in the same satellite spot beam. In this case no OBP switching/routing functionality is required in the satellite payload, and only multiplexing may need to be supported. A typical example of OBP processor operating in this scenario is SkyPlex.

Regardless of the network configuration (for star or mesh operation), in general the RCST transmits based on the DVB-RCS air interface and receives based on the DVB-S/S2 air interface, while the air interfaces used by RSGW for transmit and receive depends on the OBP specifics. In the case of RSM B systems for example, the RSGW transmits based on the DVB-RCS air interface and receives based on the DVB-S air interface, since the RSGW can be assimilated with a large RCST.

Considering the RSM-B systems in mesh configuration as an example, the following single-hop communications within an Interactive Network are supported:

- RCST-to-RCST (mesh communication):
  - the transmitting RCST uses a DVB-RCS air interface;
  - the receiving RCST receives according to a DVB-S/S2 air interface.
- RCST-to-RSGW (star communication):
  - the RCST transmits using a DVB-RCS air interface;
  - the RSGW receives according to a DVB-S/S2 air interface.

- RSGW-to-RCST (star communication):
  - the RSGW transmits according to the DVB-RCS air interface;
  - the RCST receives according to the DVB-S/S2 air interface.
- RSGW-to-RSGW:
  - the transmitting RSGW transmits according to the DVB-RCS air interface;
  - the receiving RSGW receives according to the DVB-S/S2 air interface.
- RSGW to RSGW communications within the same IN could be used to solve situations such as:
  - corporate access, providing interconnectivity between two remote networks part of the same corporation (e.g. different divisions). Each RSGW will act as traffic concentrator for one division of the corporate network (see figure 4.4);
  - multiconference scenario, allowing the presence of an MCU (Multi-Conference Unit) behind each RSGW, in order to support a multipoint to multipoint connection.

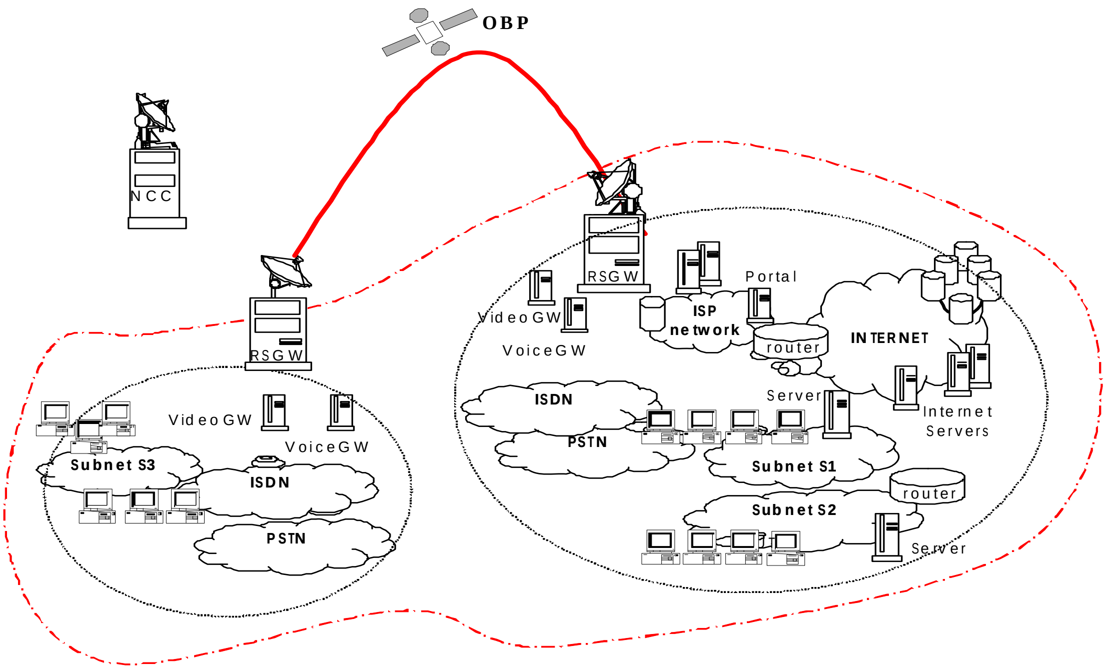

Figure 4.4: RSGW to RSGW Communication: Corporate Networks Interconnection

As in the case of transparent scenarios, if multiple INs are supported in a given single-beam satellite system, multiple NCCs are envisaged (one per IN), all located in the spot beam coverage. Each IN may have one or more RSGWs. The entire population of RCSTs and RSGWs is logically divided in separate groups, one per IN. Only communications between RCSTs and RSGWs belonging to the same IN are possible in one hop. The resources in each IN are separately managed from the resources in any other IN. The partition of satellite resources among INs is part of the centralized management and control functions of the Network Operations Centre (NOC), which is also responsible for the coordination among NCCs.

The resource of an IN can in turn be logically partitioned among different and isolated satellite sub-networks, e.g. each associated with a Service Provider. Each sub-network will be formed by a group of RCSTs and one or more RSGWs. The network can be configured in such a way that one-hop communications can only be established between RCSTs and RSGWs belonging to the same IN sub-network. Within an IN sub-network an RCST may be configured to access a predefined RSGW, based on various criteria (e.g. type of service, QoS, network congestion).

4.2.1.2 Multi Beam

In case of the multi-beam scenario, the RCSTs, the RSGW(s) and the MS can be located in different satellite spot beams. In this case, the OBP needs to implement switching/routing functions, in order to switch/route connections/packets from one uplink beam to a downlink beam, according to the final destination to be reached. Typical examples of OBPs operating in this scenario are AmerHis [i.27] and WEST OBPs [i.30].

In mesh configuration an RCST can communicate with other RCSTs, in the same or to different spot-beams. Apart from this aspect, all other considerations identified for the single-beam scenarios will also apply to the multi-beam case.

In star configuration the RCSTs, located in the same or in different spot beams, communicate with an RSGW located in a spot beam which may coincide or not with one of the beams where the RCSTs are located.

All RCSTs and RSGWs, in either mesh or star topologies, can communicate with the MS thanks to the OBP switching/routing functions.

5 Definitions of C2P Core Elements

5.1 Channels

The channels are used in the context of DVB-RCS networks for the purpose of dynamic resource control (DAMA scheduling). As such, they refer to pools of return link (uplink) resources in a user beam (timeslots in the MF-TDMA frame), identified in RCSTs and NCC by different Channel_ID values. A Channel_ID value is used by the RCST to tag the capacity requests relative to a given channel, and by the NCC (DAMA Scheduler) to tag the capacity allocations (in TBTP) relative to the same channel.

The DVB-RCS standard [i.1] allows the use of channels (and Channel_IDs) for the purpose of either connectivity or QoS differentiation, or both. The present document will therefore distinguish, where appropriate, between "connectivity channels" (clause 5.1.1) and "QoS channels". The latter will preferably be designated as QoS classes, in the sense of MAC Classes of Service (MAC CoSs) (clause 5.1.2), which are a reflection at the MAC layer of the classes of service implemented at IP layer (i.e. the DiffServ's PHBs/IntServ's CoSs). For the QoS model adopted for C2P please refer to clause 7.2.1.

Each channel, whether defined for connectivity or QoS differentiation, is uniquely identified in an uplink user beam. According to version 1.5.1 of the DVB-RCS standard [i.1], Channel_ID is coded as an 8-bit word in the SAC field, the TBTP and the Information Element (IE) used with the DULM method. This allows a maximum of 256 channels to be configured in a beam, for either connectivity or QoS differentiation. This is considered sufficient for all DVB-RCS networks envisaged in short and medium term. However, the Channel_ID can be extended to 12 bits if needed, as described in annex A.

In system designs that allow the use of Channel_ID for both connectivity handling and QoS differentiation at the same time, the NCC and RCSTs should be aware of the usage of each configured channel and its Channel_ID (for connectivity or for QoS). In general a connectivity channel can be uniquely identified by the list of Channel_IDs assigned to the MAC CoSs supported by that connectivity channel. It will be convenient to use a simple logical identifier for a Channel_ID list, for example for expedite referencing the connectivity channels within the connection control protocol (i.e. in C2P messages). Some implementations may chose to use the Route_ID for this purpose (clause 5.2), but this usage will not be mandated by C2P specifications.

From the RCST viewpoint the Channel_ID has local meaning as a unique identifier used in the interaction between that particular RCST and the NCC. The same Channel_ID value can be used in different RCSTs in reference to different QoS classes/connectivity channels from those defined in the user beam. However, at the NCC level there should be a unique identification of all channels defined in a user beam for all RCSTs, and the NCC should keep a table with the mapping of the Channel_IDs configured in all RCSTs to the Channel_IDs used in the NCC, if they are different. In practice, if the same MAC QoS in a given connectivity channel is supported in different RCSTs, it will be desirable to be tagged with the same Channel_ID.

5.1.1 Use of Channels for Connectivity

When channels are used for connectivity, they correspond to RL/UL physical partitions associated with different destination downlinks. In the most general case a partition may include slots anywhere in the MF-TDMA frame, but constraints are typically imposed in order to simplify the resource configuration and control (e.g. a partition can be defined to include a contiguous set of slots on an MF-TDMA carrier or an integer number of full MF-TDMA carriers).

A connectivity channel can be defined with reference to an RCST (RCST-centric) or to an uplink beam (beam-centric):

- if a connectivity channel is RCST-centric, it is accessible to only one RCST and the corresponding partition (in terms of capacity/rate) should not exceed the RCST transmit rate (i.e. the equivalent of a full MF-TDMA carrier);
- if a connectivity channel is beam-centric, its capacity is shared among several RCSTs and the channel capacity can exceed the RCST transmit rate, i.e. the channel partition can span two or more MF-TDMA carriers.

Connectivity channels are implicitly associated with destination downlinks. Depending on the network reference scenario, a destination downlink can be a TDM carrier or a MF-TDMA carrier group in a downlink beam. In the case of transparent mesh scenarios with terminals equipped with single carrier demodulators, a destination downlink can be reduced to a single MF-TDMA carrier. Moreover, in such scenarios the destination downlink of an RCST-centric connectivity channel could be defined even in terms of slots on a MF-TDMA carrier.

In general, multiple channels can be established from a given RCST/uplink beam to different destinations. In a mesh network for example, the channels can be used to differentiate user data from signalling: user data are channelled to user beams where the RCSTs are located or to a Gateway beam, while signalling is channelled to the beam where the NCC is located.

Connectivity channels are in general under the control of the network operator and are typically static or at best quasi-dynamic, and should be consistent with the connectivity/switching matrices implemented in various elements of the network (function of network specifics), primarily in the OBP.

When multiple channels are defined in an uplink beam, the capacity (slots) allocated by the DAMA Scheduler to an RCST in all channels should satisfy the following RCST constraints:

- the total capacity should not exceed the RCST transmission rate;
- the slots should not be allocated at the same time (same timeslot) on different MF-TDMA carriers.

The above constraints can be accounted for at two levels:

- at channel definition level, if the channels are RCST-centric. The RCST-centric channels should never include slots overlapping in time and the total number of slots in all channels should not exceed one MF-TDMA carrier worth of slots;
- at DAMA Scheduler level, especially if the channels are beam-centric and span multiple carriers. The DAMA Scheduler should not allocate slots to an RCST in the same timeslot on different carriers of the same channel or of different channels.

5.1.1.1 Connectivity Channels in the Case of Continuous Carrier Return Links

In transparent network reference scenarios with return links based on Continuous Carriers (CCs), each CC can be associated with an RCST-centric connectivity channel.

5.1.2 Use of Channels for QoS

When channels are used for QoS differentiation, they correspond to logical partitions of the RL/UL resources, associated with different QoS classes (i.e. MAC CoSs). The partitions are defined in terms of number of slots (i.e. they are scalars), and the slot position within the MF-TDMA frame structure is irrelevant. The DAMA Scheduler can allocate slots to an RCST anywhere in the MF-TDMA frame (more accurately in a connectivity channel partition, if connectivity channels are used) for any QoS class, while observing the RCST constraints.

Similar to the connectivity channels, the "QoS channels" (i.e. MAC CoSs) may also be accessible to a single RCST (RCST-centric) or to several RCSTs, all in the same uplink beam (beam-centric), depending on the network reference scenario. In the latter case the QoS channel capacity can exceed the equivalent of an MF-TDMA carrier, but this will not raise issues with the generation of the TBTP, since the channels are scalars. It is therefore easy to avoid allocating overlapping slots to an RCST for different QoS classes, since the slots can be anywhere in the MF-TDMA frame; for example, the allocation can be in the form of contiguous slots on two successive carriers, or can follow a certain QoS class-specific pattern (e.g. to satisfy the jitter requirements).

The use of multiple channels for QoS differentiation is in line with the current trend in the QoS design for DVB-RCS access networks, as defined by the SatLabs group [i.4].

Within SatLabs QoS Specifications there is no definition of Channels. However, Channel_IDs are used to identify the Request Classes (RCs), which correspond to the MAC Classes of Service defined at NCC level, as seen from the point of view of the RCST.

5.2 Route and Route_ID

The Route_ID has been originally introduced in the DVB-RCS standard to facilitate "label switching" on board of regenerative satellites of uplink packets to destination downlinks; as such it identifies a destination forward downlink. In version 1.5.1 of the DVB-RCS standard its definition has been extended to allow for other usages, e.g. to indicate a connectivity channel, which is associated with a destination downlink as well (see clause 5.1.1). A path for packet transmission from an uplink beam to a destination downlink constitutes a route.

Route_ID is defined as a16-bit subfield in the SAC field and in the IE used with the DULM method.

There is no conflict between the two usages of the Route_IDs, since the Route_IDs are not needed for both label switching and connectivity channels in the same system. On one hand, the Route_IDs are not needed for on board switching in the case of regenerative satellites based on routing/switching tables and in the case of transparent satellites, the corresponding bits in the Route_ID subfield (or a subset thereof) could be used for other purposes. In such cases and in the context of C2P, for example, they can be used in the C2P messages to identify a list of Channel_IDs used for the MAC CoSs within a connectivity channel (see clause 5.1). This information can be exploited by the RCST to optimize the filling of the slots assigned within a connectivity channel with packets from different Request Classes, consistent with the QoS requirements in [i.7]. Such usage of the Route ID is implementation specific and is completely transparent to the satellite payload (i.e. the Route_ID is not extracted/used in the OBP). On the other hand, in systems supporting label switching (based on Route_IDs) there is no need for connectivity channels, therefore the usage of Route_IDs in such systems is completely transparent to the ground equipment (e.g. RCSTs). In other words, the Route_ID is used either on board for label based switching (prefix method) or in RCSTs for QoS purposes.

5.3 IP flows

IP flows typically refer to elementary flows generated by various applications, but can also refer to aggregate flows sharing some common properties (e.g. the same service class/PHB, same source/destination, etc.). The IP flows are identified by various criteria, such as IP source/destination addresses, source/destination port number, DSCP, etc.

IP flows are carried as part of the streams associated with different connections, therefore a flow can be associated (by inheritance) with the connection's parameters (see clause 5.4). A connection can carry one or multiple IP flows.

5.4 Connections and streams

5.4.1 Connections

5.4.1.1 Definition

A connection is understood as a logical association between two entities or network reference points required for traffic transmission between two communicating parties. The reference points constitute the end-points of the connection. In the case of multicast/broadcast connections one entity is a virtual entity, representing multiple end-points.

Connections can be defined at different layers of the protocol stack (e.g. MAC connections, IP connections, TCP connections).

In the context of C2P a connection refers to a MAC connection enabling the transmission of packets in MAC format from one network reference point to another (unicast) or to multiple (multicast or broadcast) network reference points. The packets are generated by a source and are intended for a destination. The source and destination, as communicating parties, are associated with the two reference points of the connection, therefore one can refer to source/destination reference points or parties. The primary User plane connotation of the source and destination is thus extended to the Control plane and used from the C2P point of view.

For a unicast connection the reference points can be two RCSTs or an RCST and the NCC/GW (or RSGW). For a multicast connection one reference point (the multicast source) is typically the NCC/GW (or RSGW) and the other reference points correspond to other RCSTs. The source reference point can also be an RCST, in which case the other reference points correspond to other RCSTs and/or GWs. Since the connection is defined at MAC layer, the reference points are identified by their MAC addresses. Each connection is associated with a given QoS class, designated as C2P Class of Service (C2P CoS). An additional identifier is needed for the purpose of encapsulation/reassembly at the reference points of the connection. The VPI/VCI (ATM profile) or PID (MPEG profile) can be used to this end.

In summary, a connection is defined by the following attributes:

- a pair of RCST MAC addresses (source and destination);
- C2P Class of Service (C2P CoS);
- VPI/VCI pair (ATM profile) or a PID (MPEG profile), to be used for encapsulation/re-assembly at the connection reference points.

NOTE 1: In the case of MPEG2-TS format the destination MAC address is included in the DSM-CC header during the encapsulation process.

NOTE 2: For multicast connections, the destination MAC address for each end-point of the connection is derived from the IP multicast address.

NOTE 3: For connections triggered by an incoming IP packet (see clause 7.1), the destination IP subnet of that packet may also be considered as an attribute of the connection.

NOTE 4: The VPI/VCI pair is alternatively referred to as VCC.

The above attributes can be "lumped" under a unique logical identifier (Connection Reference ID) for the purpose of connection management by the NCC (e.g. for the referencing of a connection in the C2P messages relative to that connection). The Connection Reference ID can be set by the RCST (RSGW) or by the NCC, to values (coded as 2 bytes) taken from two disjoint pools, as shown in figure 5.1.

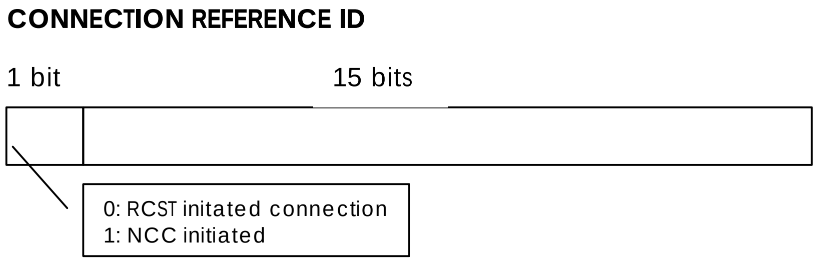

Figure 5.1: Connection Reference Identifier

When set by the NCC, the Connection Reference ID is unique in the entire network controlled by the NCC. When set by the RCST, the Connection Reference ID has local meaning, allowing the RCST to locally identify an active connection.

Due to its local nature, the RCST-defined Connection Reference ID assumes a globally valid uniqueness only if used in conjunction with an RCST identifier (e.g. the MAC address). One can therefore conclude that a connection can be uniquely identified at the network level by:

- a Connection Reference ID set by the NCC, in the case of NCC-initiated connections, **or**
- a Connection Reference ID set by the RCST, plus an RCST identifier (the MAC address), in the case of RCST-initiated connections.

NOTE 1: The Connection Reference ID set by the NCC will be used in all C2P messages relative to that connection.

NOTE 2: In the case of an RCST-initiated connection between two RCSTs, the connection Reference ID set by the RCST is used in the C2P messages exchanged between the initiating RCST and NCC, while a Connection Reference ID set by the NCC can be used in the C2P messages exchanged between NCC and the other RCST.

The above definition of the connection attributes provides a unidirectional view for the connection. The connection attributes are linked to the transmission satellite interface of the RCST. For a bidirectional or duplex connection each direction will correspond, from C2P point of view, to a unidirectional connection. A duplex connection will therefore be identified in general by a pair of Connection Reference IDs, as illustrated in figure 5.2 and figure 5.3 for star transparent and mesh regenerative systems, respectively. In some special cases (see clause 5.4.3.5) the same Connection Reference ID can be used for both directions, i.e. for the two unidirectional connections.

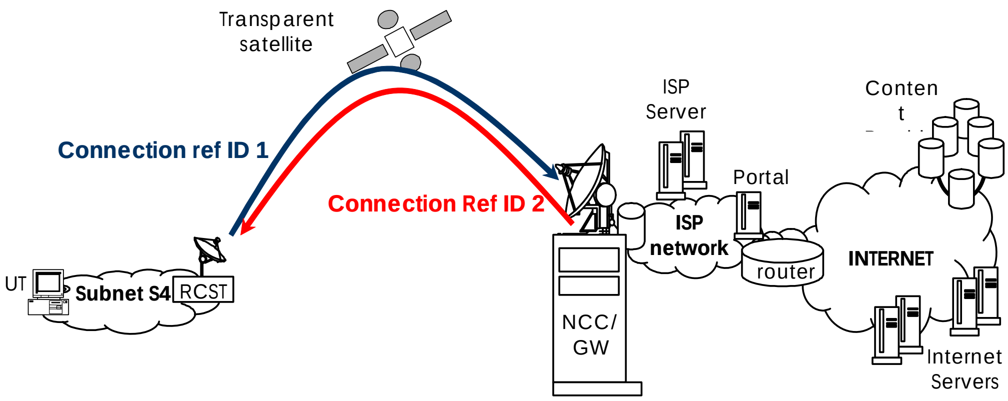

Figure 5.2: Star Connections in DVB-RCS Transparent Systems

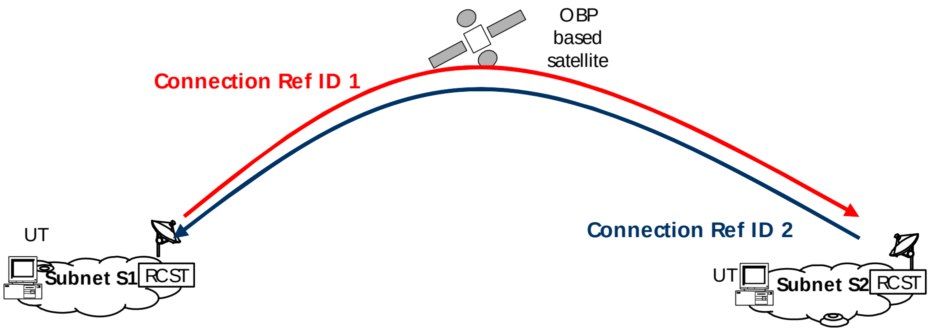

Figure 5.3: Mesh Connections in DVB-RCS Regenerative Systems

Multiple connections can be established between two reference points, e.g. for different classes of service.

EXAMPLE: An RCST can establish multiple connections with the same destination (RCST or NCC/GW), associated with different C2P CoSs.

In general, the number of distinct connections that can be established between any two network elements (RCSTs, NCC/GW) is given by the maximum number of C2P CoSs implemented in the system. For each connection a QoS profile will be defined, consistent with the connection QoS class (i.e. C2P CoS).

Since a connection established from a given terminal with a given C2P CoS/QoS profile is always mapped to a MAC CoS within a connectivity channel, the connection can be associated with a number of additional parameters, namely:

- (Source) Terminal Group_ID/Logon_ID.
- Channel_ID of the MAC CoS.
- List of Channel_IDs for the connectivity channel associated with a destination downlink.

NOTE: The Route_ID can be used in C2P messages for easy referencing a list of Channel_IDs.

- QoS Profile.

All the above parameters and attributes of a connection are considered layer 2 parameters. With the exception of Group_ID and Logon_ID (which are assigned per session at RCST's logon), they are assigned each time the connection is established, for the duration of the connection, and are configured in the relevant network components (RCST/RSGW, NCC/GW). The assignment can be done either by configuration or dynamically, by using the Connection Control Protocol (C2P). The connection modify command can then be used to change the QoS Profile.

The various types of connections that can be established are described in clause 5.3.3.

5.4.1.2 The Class of Service of a Connection

As already indicated, a connection with a given QoS class (i.e. C2P CoS) is customarily mapped to a MAC Class of Service (MAC CoS) supported by the NCC (i.e. by the DAMA Scheduler in NCC). A MAC CoS is seen from the RCST's point of view as a Request Class (RC), according to the SatLabs' QoS model [i.4].

SatLabs QoS model follows the DiffServ [i.14] framework. Other QoS models, such as that defined by ETSI BSM [i.3], are more general in scope and accommodate both the DiffServ framework and the IntServ [i.16] framework; the BSM model also addresses end-to-end and upper layer QoS issues.

Both DiffServ and IntServ architectures rely on IP classes of service, but their definition is different:

- In DiffServ the classes of service are defined for IP flow aggregates and have local (per hop) meaning. They are referred to as Per Hop Behaviours (PHBs) (e.g. EF, AF, BE PHBs) and are mapped to C2P CoSs.

- In IntServ the classes of service (Guaranteed Service - GS and Controlled Load -CL classes) are defined per IP flow and have an end-to-end meaning. They correspond to end-to-end IP connections set-up hop-by-hop, typically by using the RSVP [i.19] protocol (or NSIS [i.23] protocol or SIP [i.20] proxy), which allows to dynamically establish the parameters of the connections and configure them in all network nodes along the end-to-end path.

In the satellite network (seen as a node on the end-to-end path) each IP connection could in principle be mapped to a C2P connection, thus enabling higher granularity (at IP flow level) with regard to QoS differentiation at the link layer, when compared to DiffServ. However, this raises scalability issues since the number of IP flows in a network can be rather large (and it is prohibitive to keep per-flow states in all network's components), therefore several IP connections (and the corresponding IP flows) with similar service requirements are typically aggregated (mapped) into a C2P connection with adequate C2P CoS. All C2P connections within a connectivity channel with the same C2P CoS will then be aggregated into one RC/MAC CoS and will share the bandwidth available for that MAC CoS. The need for IP flow aggregation was also recognized in the BSM QoS model (see figure 7.2 in clause 7.2.1.1, where the IntServ IP flows were scheduled into a reduced number of queues at the SI-SAP interface (one queue for GS class and another queue for CL class), each identified by a Queue ID (QID). Two schedulers are thus envisaged at the IP layer (part of Satellite Independent Access Function - SIAF), one for each IntServ CoS (i.e. GS and CL). The scheduling discipline should in general provide fairness among the IP flows, while allowing priorities between them, if needed (e.g. it can use a weighted fair queuing algorithm).

Priority scheduling can also be applied in the case of DiffServ model, for example to schedule packets corresponding to the same application (same PHB) but with different priorities, into C2P connections with the same C2P QoS, to be aggregated into a corresponding MAC CoS. The priorities can be established based on filter patterns, including the DSCP as a primary identifier/differentiator. The differentiations between IP flows within the same PHB will require additional queues at IP level, e.g. one queue per DSCP.

NOTE: With no priorities the DiffServ model only requires one IP queue per PHB, unlikely the IntServ model which always require one IP queue per IP flow.

From the above description one concludes that within both DiffServ and IntServ frameworks, the IP flows or flow aggregates are mapped to connections with defined C2P CoS, each corresponding to an RC/MAC CoS. C2P CoSs are thus consistent with both the SatLabs QoS model and the BSM QoS model. As a result of C2P CoS to MAC CoS mapping, the DAMA Scheduler at NCC will be capable of providing adequate QoS support for connections.

5.4.1.3 Connections in the Context of Continuous Carriers Return Links

The definition of connections in clause 5.4.1 also applies in the case of return links based on a Continuous Carrier (CC). The class of service of a connection originating at an RCST can still be based on the DiffServ or IntServ architectures, but the overall QoS model is slightly modified, resembling the QoS model used on the forward link in NCC/GW, but with a much lesser degree of traffic aggregation (limited to RCST's traffic). Since a CC is fully assigned to the RCST for a determined period of time (e.g. for the duration of the logon session), there is no need for dynamic resource requests and allocations (i.e. DAMA mechanism). The RCs in the SatLabs QoS model are therefore not applicable (not needed). The CC bandwidth will be simply partitioned between connections, while trying to satisfy the service requirements of each admitted connection (consistent with the connections C2P CoS). These partitions are similar to the assignments for different MAC CoSs, but are fixed for the duration of the connection (i.e. they are setup at connection establishment ant not subject to changes to track traffic variability). In this context one can still claim that a C2P CoS is mapped to a MAC CoS, seen as a "MAC pipe". The MAC pipes may have hard or soft boundaries, to reflect various C2P CoS requirements. Policies could be defined and configured in RCSTs, in order to allow resources sharing among connections. All these are MAC mechanisms, equivalent to the DAMA mechanisms used for standard MF-TDMA based return links. The IP mechanisms are the same, regardless of the nature of the return link.

5.4.1.4 Connections in the Context of Regenerative Satellite Systems

In the case of regenerative satellite systems, the encapsulation format on the uplink may be subject to changes in the OBP (depending on the OBP architecture), either to match the format used on the downlink or to match the format used for switching/routing, or both. This may require the use of the VPI/VCI pair or PID on-board the satellite, for the purpose of reassembly/encapsulation. As an example, if the OBP switch is implemented as an IP router, the IP datagrams will have to be reassembled before routing and then segmented/encapsulated in the format defined for the downlink. In this case the end-to-end connection could be split into an uplink connection and a downlink connection, with the satellite becoming a reference point for both connections. From the C2P point of view the satellite would appear as a network node and should be involved in the establishment and termination of connections. Such scenario is rather unlikely, at least in the short term, since most of the envisaged OBP architectures rely on table-based switching and not on true packet switching. With such architectures the continuity of a connection on the downlink is automatically guaranteed and there is no need to separately manage the downlink part of the connection. Furthermore, the extension of the DVB-RCS standard to future regenerative satellites (RSAT) [i.2] provides an efficient means for "label-based switching" of packets at the link layer, based on the use of the Route_ID, with or without uplink-downlink format change. The end-to-end connection definition can therefore be maintained, provided that adequate congestion control schemes are devised in order to provide high probability of resource availability on the downlink part of the connection, matching those on the uplink part of the connection; such schemes will depend, among others, on the approach to QoS and downlink resource management in the OBP. In conclusion, it is felt that it is not necessary at this time to consider the satellite as a reference point for the definition of connections. Therefore in the current revision of the present document (and for the purpose of C2P specification) the connections will be considered between on-ground network reference points, as defined in this clause.

5.4.2 Streams

In the context of C2P a stream is interpreted as the logical flow of layer 2 data packets from one network reference point (e.g. RCST, NCC/GW) into the satellite/DVB-RCS network. The notion of stream, as well as that of connection, is linked to the transmission satellite interface. The stream flow is in the form of MAC packets resulting from the segmentation and encapsulation of IP datagrams. PIDs or VPIs/VCIs are used for the purpose of segmentation/encapsulation and reassembly at the reference points of the connection, therefore they are seen as parameters of the connection and also of the associated stream.

A stream is uniquely identified (for segmentation, encapsulation and reassembly operations) by a Stream_ID - as logical identifier, and all data packets in the stream share the same Stream_ID. For the MPEG profile the Stream_ID corresponds to the pair {PID, Destination MAC address}. For the ATM profile the Stream_ID corresponds to a {VPI/VCI} pair.

The concept of stream may be different in transparent and regenerative systems:

- In transparent systems (figure 5.4), the stream can be seen as the logical transmission of data packets from one network reference point to another (unicast) or to a group of network reference points (multicast). All reference points are on the ground and correspond to those used for the definition of the underlying connection. The stream is thus associated with a connection. The satellite is transparent to the stream transmission, i.e. the Stream_ID is not used on board the satellite for any processing. The two streams represented in figure 5.4 (Stream 1 and Stream 2) correspond to the two directions of a bidirectional unicast connection.

- For OBP based systems (figure 5.5), the concept of stream considers the routing /switching capabilities of the OBP (e.g. the OBP may switch based on the Stream_ID - either a PID value or a VPI/VCI pair). The stream will be defined from the network reference point on ground to the reception interface of the OBP. This definition allows for the potential format change in OBP, if required by the design. For these systems the end-to-end connection is split into an uplink connection and a downlink connection, both having the satellite as one of their reference point (see comments in clause 5.4.1.4). The stream, as defined in this clause, will be associated with the uplink connection (Stream 1 in figure 5.5). Two streams (Stream 2 and Stream 3) are associated with downlink connections to different RCSTs in the same downlink beam, as illustrated in figure 5.5, which applies to a multicast scenario. A downlink stream may include aggregates of packets from multiple uplink streams. This is in general true whenever there is a multiplexing function at the OBP (e.g. as a result of switching), whether the format is changed or not. In the case of OBPs based on label switching it is expected that the downlink connections are dimensioned to carry the packets in all aggregated downlink streams, or that effective congestion control schemes are used. In the case of OBPs relying on table based routing (e.g. RSM-B systems), there is no on-board congestion and the aggregation of streams at OBP is equivalent to the aggregation of connections, defined with regard to on-ground network reference points (as for transparent systems).

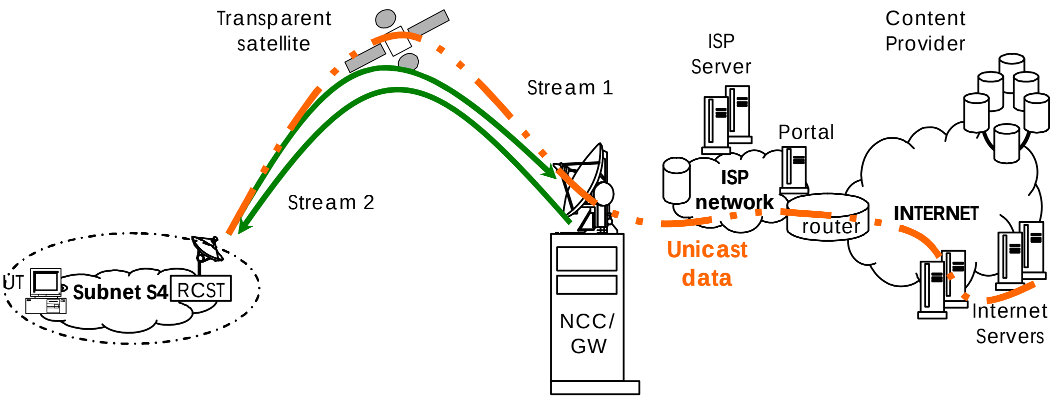

Figure 5.4: Streams in Transparent DVB-RCS Systems (Bidirectional Unicast Illustrated)

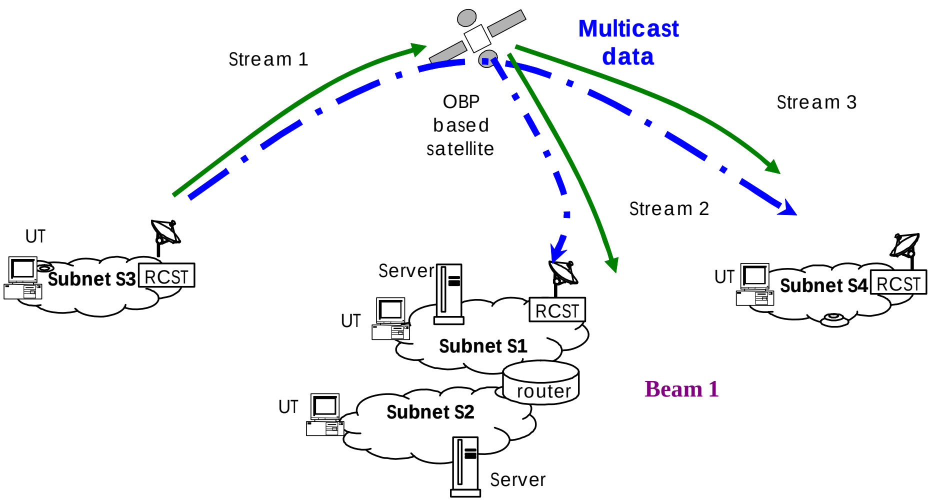

Figure 5.5: Streams in Regenerative DVB-RCS Systems (Multicast Illustrated)

A stream can carry one or several IP flows (i.e. aggregation of IP flows, as in DiffServ model). Moreover, the notion of stream can be extended to include aggregation of streams at RCST level, corresponding to several connections within the same connectivity channel, under the same {VPI/VCI} or {PID, destination MAC address}. As a result of aggregation, the packets from different connections will be interleaved. This will have impact on packet delivery (QoS), therefore it is preferable to restrict the aggregation to connections with identical/similar QoS requirements (C2P QoS Classes mapped into the same RC/MAC CoS, identified by the same Channel_ID).

5.4.3 Connection Types

In general the connections are defined for star and mesh topologies, for traffic or signalling. Furthermore, the traffic connections can be set-up either by the NCC (NCC-initiated connections) or by the RCST (RCST-initiated connections) and can be of various types: point-to-point unidirectional or bidirectional, point-to-multipoint, star or mesh. A bidirectional connection assumes two unidirectional connections.

The various categories of connections are reviewed in the following clauses.

5.4.3.1 Star and Mesh Connections

Connections are differentiated in star and mesh, depending on whether they involve an NCC/GW or not:

- a star connection is established between an RCST and NCC/GW (RSGW in a regenerative system), i.e. there is an NCC/GW (or RSGW) involved in the communication;
- a mesh connection is established between two RCSTs, and represents direct communication between two RCSTs without the involvement of the Gateway.

NOTE: A connection between two RSGWs could be considered star or mesh, depending on the context.

This differentiation applies to both transparent and regenerative satellites that can be configured to support mesh and star topologies. In general, a network supporting mesh connections also supports star connections simultaneously (e.g. mesh overlay), primarily for signalling (to/from NCC) but also for access traffic (to/from Gateway).

In the case of star networks all traffic connections are of star type.

In the case of mesh networks the traffic connections are in general of mesh type (between RCSTs), but can also be of star type (between RCSTs and GW).

5.4.3.2 Signalling and Traffic Connections

The differentiation between

- signalling connections, and
- traffic connections

is based on the type of data conveyed by each connection.

Signalling connections carry various signalling messages (e.g. protocol-related messages, management and control messages).

Traffic connections carry user traffic.

Signalling connections could be mixed with traffic connections in the same connectivity channel and/or QoS class.

5.4.3.2.1 Signalling Connections

All C2P-based connection oriented DVB-RCS systems will require signalling connections to carry signalling messages such as C2P messages (in DULM/TIMu formats), management and control messages (e.g. SNMP messages) and other messages (e.g. for application session establishment, QoS management, etc.).

Signalling messages will be conveyed between RCSTs and NCC/GW (in a transparent system) or NCC/NMC (in a regenerative system).

Any DVB-RCS system supporting C2P will require at least one signalling connection, to be set-up just after logon, for C2P messages. It is recommended to also have a separate signalling connection for management messages. (e.g. SNMP messages). Additional signalling connections can be defined, as required by the end-to-end system design, for application session signalling, QoS signalling, etc.

For signalling connections implicitly opened at terminal logon (i.e. without the need of C2P messages), all the information required for connection establishment (e.g. transmission and reception PIDs or VPI/VCI values, MAC addresses, or IP addresses) is contained in the logon messages received by the RCST.

5.4.3.2.2 Traffic Connections

Traffic connections are established to convey user data. The connections are set-up with pre-defined profiles (service level) and therefore involve a Connection Admission Control (CAC) function. They can be classified according to different criteria:

- depending on the network entity that requests the set-up of the connection, one can distinguish between two types of connections:
  - NCC-initiated;
  - RCST-initiated.

The establishment/release of NCC-initiated or RCST-initiated connections relies on the C2P protocol, which is triggered by well-defined events at the NCC/GW and/or RCST/RSGW (e.g. arrival of an application packet, interception of an application session initiation message, NCC decision, etc).

- With regard to the type of casting, the (MAC) connections can be classified as unicast or multicast:
  - a unicast connection is a point-to-point connection;
  - a multicast connection is a point-to-multipoint connection.
- A unicast connection can further be classified in one of the following types:
  - unidirectional connection;
  - bidirectional (or duplex) connection.

All the above types of traffic connections will be briefly described in the following clauses.

5.4.3.3 NCC Initiated and RCST Initiated Traffic Connections

NCC-initiated connections are set-up/scheduled for those RCSTs that are already logged-on into the system. They can be of either star or mesh type:

- a star NCC-initiated connection is established between an RCST and a GW via a transparent satellite, or between an RCST and an RSGW via a regenerative satellite;
- a mesh NCC-initiated connection is established between two RCSTs or two RSGWs.

Each connections, whether initiated by NCC or RCST, has logical resources allocated at connection establishment for the duration of the connection, and some "entry level" or initial bandwidth resources. The latter can be altered/renegotiated during the life of the connection (e.g. by relying on the "connection modify" command of C2P), in response to changing traffic profile or as a result of new applications becoming active.

All NCC-initiated connections are fully controlled by the NCC, i.e. they are setup and released by the NCC/GW and cannot be released at RCST's request.

The RCST-initiated connections are established and released dynamically after the RCST has been logged-on to the network, upon explicit requests from RCSTs. They may also be released upon explicit request from the NCC, normally to solve a corner case.

NOTE 1: A corner case is defined as an unresolved or unstable situation of the C2P protocol due to an unexpected exchange/behaviour of messages, such as messages lost, cross request of connections (the same connection being requested at the same time from both sides), etc.

NOTE 2: From the point of view of C2P procedures the connections initiated by an RSGW fall in the category of RCST-initiated connections.

For some RCST-initiated connections, required for example for short urgent messages, the corresponding traffic can be rather intensive, though of short duration. Accepting such connections in a fully loaded network may require the pre-emption of network resources already booked for other connections, thus affecting their services. Policies can be defined in order to prevent the pre-emption of connections for real-time applications.

5.4.3.4 Unicast and Multicast Traffic Connections

This connection classification is with regard to the type of casting:

- A unicast connection (i.e. a point-to-point connection) can be established between an RCST and an NCC/GW (transparent scenario), between an RCST and an RSGW (regenerative scenario), between two RCSTs (both transparent and regenerative scenarios), or between two RSGWs (regenerative scenario). Unicast connections may therefore be star connections (when accessing external networks via GW) or mesh connections (when providing connectivity between a pair of RCSTs/RSGWs). Unicast connections can be unidirectional or bidirectional;

- A multicast connection (i.e. a point-to-multipoint connection), can be established from the NCC/GW towards many RCSTs (star transparent scenario), from an RCST towards many RCSTs (mesh transparent/regenerative scenarios) or from an RSGW towards many RCSTs (star regenerative scenario). Multicast connections are always unidirectional.

5.4.3.5 Unidirectional and Bidirectional (Duplex) Traffic Connections

In clause 5.4.1 the connections and their associated attributes have been defined in reference to unidirectional connections and it was stated that a bidirectional connection corresponds to two unidirectional connections (in opposite directions). The majority of the attributes of a unidirectional connection (with the exception of destination RCST MAC address) are transmission parameters; they become reception parameters for the unidirectional connection in the opposite direction.

In the case of a bidirectional (duplex) connection, the RCST is awarded of the transmission and reception parameters for the two unidirectional connections. It is expected that these parameters are given in a **single set of C2P messages exchange.** This is possible if both directions of the connection share the same C2P CoS, which is the most likely case, provided that the same C2P classes of service are implemented/supported in both directions (typical of mesh bidirectional connections).

In the case where C2P CoSs in the two directions are different (e.g. for star bidirectional connections), two unidirectional connections will need to be setup, by two different sets of C2P messages exchange. This is no longer a duplex or bidirectional connection, but rather corresponds to two independent unidirectional connections.

In general each direction of a duplex connection will have a connection reference ID, that identifies the corresponding unidirectional connection and its associated transmission parameters. These will be linked to some reception parameters, corresponding to the transmission parameters of the other side (i.e. for the other unidirectional connection). The situation is illustrated in figure 5.6 for a bidirectional mesh connection.

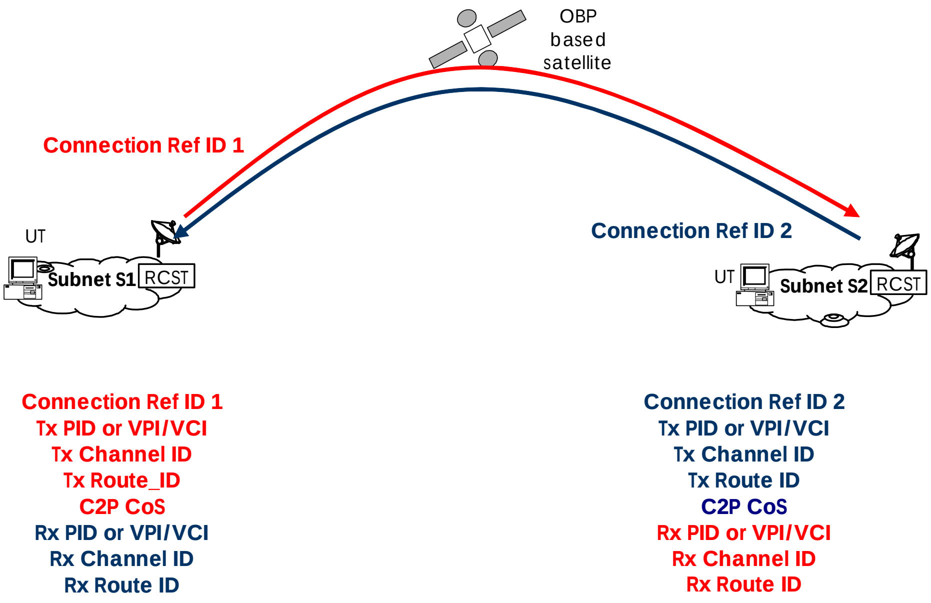

Figure 5.6: Bidirectional or Duplex Connection

The following comments can be made:

- If only one set of C2P message is required for the establishment of a duplex connection, the two sides involved in the connection will only require one Connection Reference ID - one index in the Active Connection Table (see clause 7) to fully identify and extract the parameters of the duplex connection.
- If two different sets of C2P message exchange are required, the two sides involved in of the connection will have two different Connection Reference IDs (each with its entry in the Active Connection Table), one for each unidirectional connection.

Connection Reference ID is used not only as local identifier of the connection, but also as an identifier of all C2P protocol messages associated with that connection. Whenever a terminal sends a request relative to a connection, it needs an identifier for that connection in order to recognize and differentiate the response to that request from other C2P messages it may receive.

5.4.4 C2P Connection profile

The concept of connection over DVB-RCS systems has been introduced in the previous clauses. The connection definition in these clauses is not necessarily linked to C2P signalling; it is applicable to any DVB-RCS system, supporting or not C2P. If C2P is not used, there should be other means for managing connections (e.g. by configuration under NCC control).

Connection profiles can be defined for all connections, whether C2P is used or not. If C2P is used, the profiles will be referred to as C2P connection profiles.

C2P connection profile includes all the necessary information required to establish and/or modify a connection. The information is in the form of a set of parameters/attributes linked to the type of connection, life of the connection, IP flow classification criteria, connection class of service (C2P CoS), QoS profile (consistent with the C2P CoS) in terms of transmission/reception bandwidth. This information is stored in RCST in the Connection Profile Mapping Table (see clause 7), which includes one entry (one set of parameters) for each active connection. The C2P profile parameters will be configured by management and updated via C2P messages in each RCST for all connections established from that RCST.

The C2P connection profile should include as a minimum the following information (per C2P profile entry):

- C2P profile Index;
- type of C2P connection:
  - point-to-point or point-to-multipoint;
  - unidirectional or bidirectional;
  - RCST-initiated or NCC-initiated.
- C2P CoS (corresponding to a MAC CoS);
- QoS Profile, consistent with the C2P CoS, defined in terms of connection transmission and reception bandwidth parameters (rates), such as:
  - C2P SDR/PDR return (transmission) parameters;
  - C2P SDR/PDR forward (reception) parameters.
- C2P inactivity timeout (connection inactivity timeout is used to trigger the connection release if there is no traffic).

For more details on the usage of C2P connection profile parameters please refer to clause 7.2.

6 Channel/Connection/Flow Models

The models defined in this clause, referred to as C2P models, capture in graphical form the arrangements and interdependencies between the C2P core elements, i.e. channels, connections/streams and IP flows. A general model will be defined first, applicable to all network reference scenarios. Simpler models will then be derived by particularization of the general model for specific network reference scenarios.

The consistency of the C2P models with existing QoS models and with the dynamic resource control (DAMA scheduling) will be highlighted where appropriate.

6.1 General Model

6.1.1 General Model for Beam-Centric Channels

The most general channel/connection/flow model is illustrated in figure 6.1. It applies to multi-beam network architectures with transparent or regenerative payloads, configurable to operate in star mode or mesh mode. The model is represented from the point of view of an uplink user beam (Beam A or a transponder assigned to Beam A) and is based on beam-centric connectivity channels. The general model for RCST-centric channels is described in clause 6.1.2.

A review of the key features of the model is provided below. The various logical identifiers used for the C2P entities and other relevant entities are defined at the top of the figure.

The complexity of the general model has been driven by the need to suit the mesh topologies, with either transparent or regenerative payloads, in which RCSTs in a given beam communicate with RCSTs in the same beam or other beams. To this end mesh connectivity channels need to be established between beams. A typical example of applicability of the general model is the RSM-B family of satellite networks, based on a regenerative payload (see clause 6.2.1).

In a mesh network the RCSTs also communicate with the NCC for signalling (e.g. M&C signalling, including C2P signalling), therefore connectivity channels also need to be established from each user beam to the beam where the NCC is located. Such channels are in general different from the mesh channels; they are star channels. For simplicity, all star channels used for signalling (and the connections inside) will be referred to as M&C channels/connections, even if they may be used for other signalling messages (e.g. applications session signalling, QoS signalling). Star channels can also be established for star traffic to NCC/GW (in the case of transparent scenarios) or to RSGW (in the case of regenerative scenarios). Since in the case of transparent scenarios the GW is typically collocated with the NCC, the same star channels can be shared by the signalling and traffic connections. In the case of regenerative scenarios of RSM-B type, the star channels are identical to the mesh channels, since the NCC and RSGW can be seen (from the point of view of the RF equipment) as conventional terminals with extended functionality.

With the addition of star channels, the model in figure 6.1 is applicable to multi-beam star topologies as well; the model is thus general.

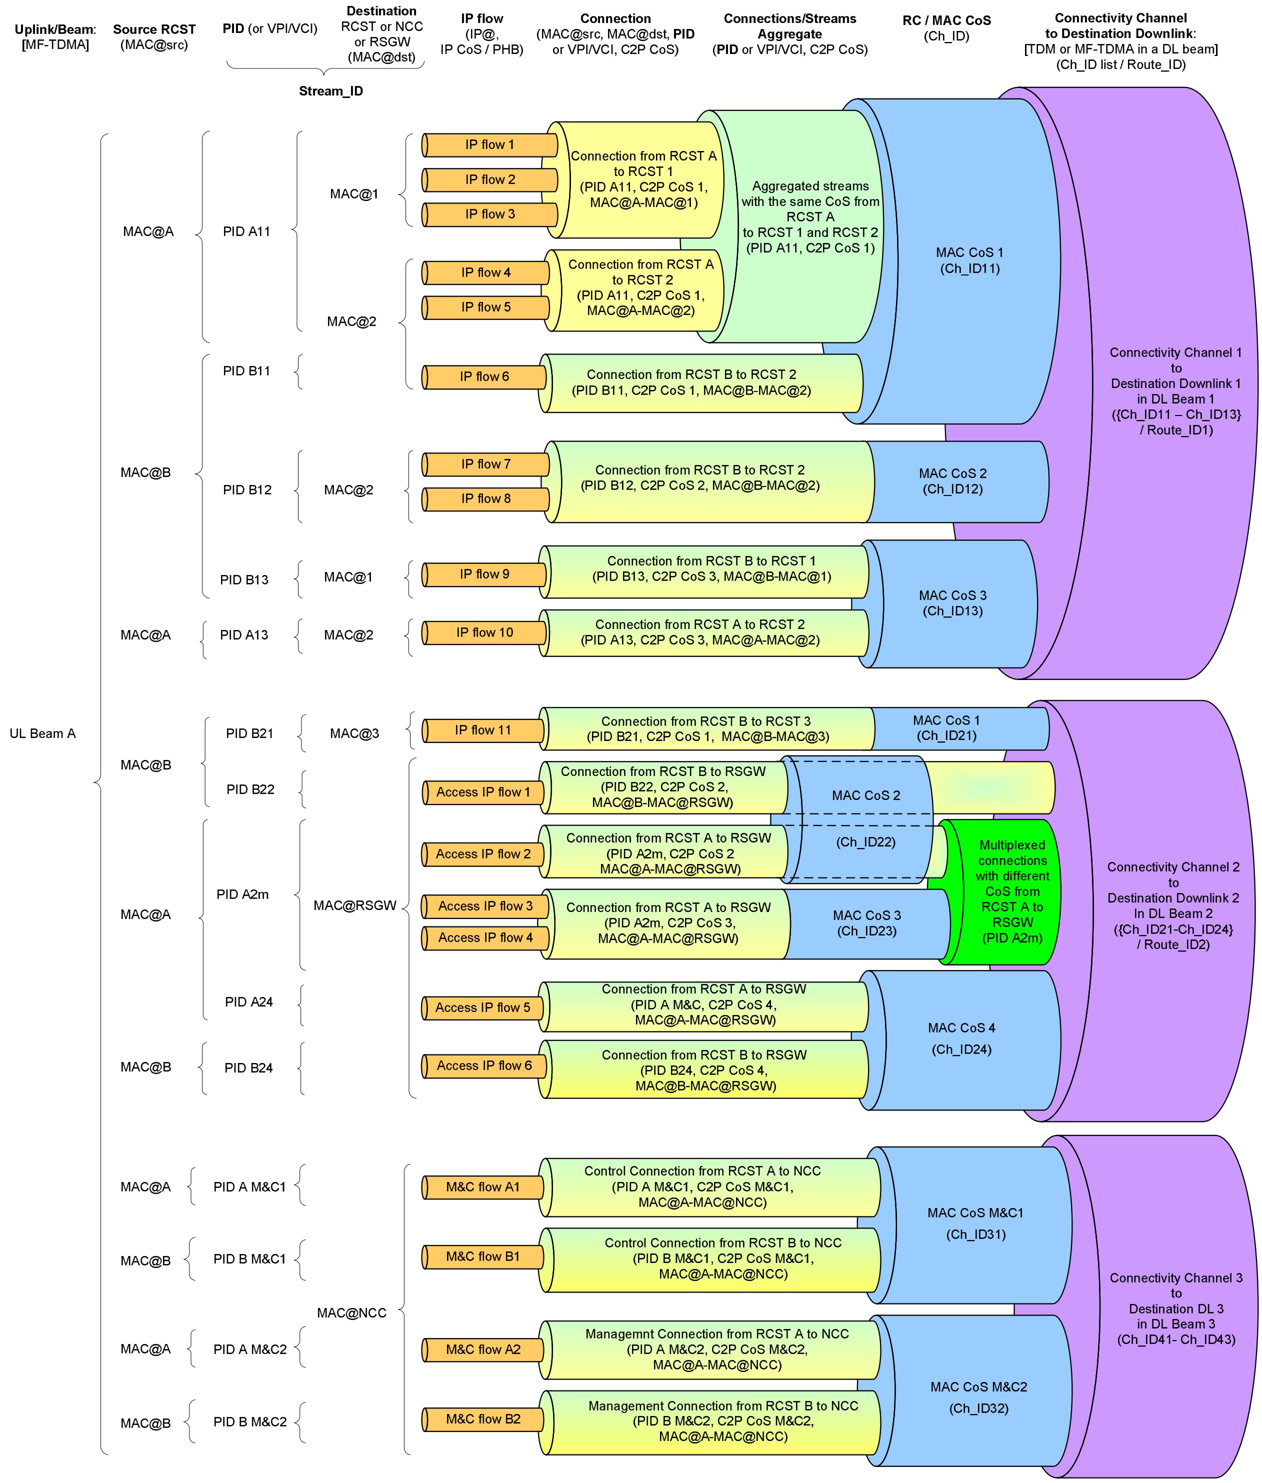

Figure 6.1: General Channel/Connection/Flow Model: Beam-Centric Channels (MPEG Profile Illustrated)

Terminals in beam A (two illustrated - RCST A and RCST B) can establish mesh connections with terminals that can be reached via different destination downlinks, e.g. RCST 1 and RCST 2 can be reached via destination downlink 1, while RCST 3 and RSGW can be reached via destination downlink 2. The downlinks rely on either TDM or MF-TDMA air interfaces, depending on the network reference scenario. Furthermore, the terminals in Beam A can also establish star connections for M&C signalling and traffic to the NCC/GW, which can be reached via destination downlink 3.

In order to establish the above connections, three connectivity channels are needed and each channel is accessible to all terminals in the uplink beam A. The three connectivity channels have been assumed in different DL beams

(DL Beam 1, 2 and 3, respectively), but more than one connectivity channel can be defined in a DL beam (e.g. corresponding to different TDM carriers).

A connectivity channel will be identified by a Channel_ID list, which could be referenced (for convenience, in the C2P messages) by a unique logical identifier - the Route_ID, as described in clause 5.2. The Route_ID used in this way will point to a destination downlink, but will not be used for on-board label-based switching.

Multiple C2P classes of service can be supported in each connectivity channel. For illustration purposes, four classes of service have been considered in figure 6.1 (C2P CoS 1, C2P CoS 2, C2P CoS 3 and C2P CoS 4), that can be implemented in all connectivity channels. An additional class of service has been considered for the destination downlink 3 (to NCC/GW), for M&C messages (C2P CoS M&C). In general the classes of service for different destination downlinks can be all different.

A terminal can establish multiple connections to the same destination (RCST, RSGW or NCC/GW), associated with different C2P classes of service. The C2P CoS of a connection will be mapped into a MAC class of service (MAC CoS).

At RCST level the MAC classes of service are also designated as Request Classes (RCs), according to SatLabs terminology [i.4]. Each RC/MAC CoS is identified by a Channel_ID. As already stated, the Channel_IDs are in general RCST-specific (local), meaning that different Channel_IDs can be used in different RCSTs for the same MAC CoS (this is a system implementation/configuration issue). However, at NCC level a given MAC CoS in a connectivity channel should be uniquely identified; a second set of channel identifiers could be used to this end, e.g. Channel_ID_NCCs. Nevertheless each Channel_ID used at RCST level is mapped to a unique Channel_ID_NCC; the mapping is RCST-specific.

If possible, it would be preferable to use in all RCSTs in an UL beam the same (unique) Channel_ID for a given RC/MAC CoS in all connectivity channel established from that UL beam. This is the case illustrated in figure 6.1, where Channel_IDxy refers to the MAC CoS y in the connectivity channel x. In this case the mapping of the Channel_IDs to Channel_ID_NCCs would be the same for all RCSTs, meaning that a unique set of Channel_IDs could be used at RCSTs and NCC (i.e. there is no need for the Channel_ID_NCCs).

The PIDs (corresponding to the MPEG profile) have been used in figure 6.1 for illustration purposes, in association with specific connections/streams, but the VPI/VCI (corresponding to the ATM profile) could have equally been used. Normally, separate PIDs (or VPI/VCIs) are used for each connection/stream originating at a given RCST (for encapsulation purposes). The situation may be different if the aggregation of connections/streams (to the same destination downlink) under the same PID (or VPI/VCI) is allowed in RCST. Two cases can be identified:

- Aggregation of connections/streams with the same C2P CoS into one MAC CoS, for the purpose of reducing the number of MAC classes of service at NCC. This is illustrated in figure 6.1 by the aggregation of streams from RCST A to RCST 1 and RCST 2 in the destination downlink 1, under the unique PID A11; the PID needs to be configured in both destination RCSTs for the purpose of packet re-assembly/filtering.

- Aggregation of connections/streams with different C2P CoSs. This is illustrated by the connections/streams aggregate from RCST A to RSGW in the destination downlink 2. The two connections in the aggregate, one with C2P CoS 2 and the other with C2P CoS3, will be identified by separate Channel_IDs (Ch_ID22 and Ch_ID23, respectively), so that MAC CoS of each connection is preserved. However, the PID (PID A2m) is shared by the two connections, therefore this type of aggregation is more a form of connection multiplexing (hence the suffix m in PIDA2m); its purpose is to reduce the number of streams (and of the logical identifiers) in the system.

The two types of aggregation are in general applicable to all destination downlinks.

Connections aggregation is actually an aggregation of streams associated with the connections, therefore the same logical identifier are used for both streams aggregation and connections aggregation (PID or VPI/VCI). When streams are aggregated under a unique PID (or VPI/VCI), the packet interleaving has to be done at IP level, and not at MAC level, otherwise the original IP datagrams could not be reassembled at destination.

In the case of mesh networks using the ATM profile, the aggregation under a unique VPI/VCI of connections/streams to different RCSTs (at the same destination downlink) may lead to the situation where an RCST could extract data intended for other RCSTs, since there is no destination MAC address in the encapsulation header. This is undesirable in terms of additional processing/filtering and may raise security concerns. It is true that the re-assembled IP packets will most probably be discarded if the destination IP@ is not part of the RCST directly connected (or routed) subnets; they will therefore not be transmitted over the RCST LAN interface. If the security or other concerns still persist, the connections aggregation under the same VPI/VCI should be limited to connections ending at the same destination RCST.

NOTE: In the case of star networks the destination is always unique (GW/NCC), therefore the aggregation under a unique VPI/VCI does no raise any concern.

It is noted that the above issue does not exist in the case of MPEG profile, for which the connections/streams are implicitly identified not only by the PID value but also by the destination MPE@, which is the MAC@ of the destination network reference point (RCST, RSGW or GW/NCC), automatically inserted in the stream during encapsulation (in the header of the DSM-CC section).

Finally, the IP flows in figure 6.1 are classified and aggregated (possibly with different priorities) into a reduced number of IP CoSs/PHBs, which are then mapped to connections with the same or different C2P CoSs. The IP CoS is a more generic term and applies to both IntServ and DiffServ architectures, while the PHB is DiffServ specific (see clause 5.4.1.2 for a discussion of the QoS models that can be accommodated by the C2P). The IP flows associated with a connection with given C2P CoS may correspond to different sessions of the same or similar user application(s). As an example, the three IP streams in the connection from RCST A to RCST 1 (at destination downlink 1) associated with CoS 1 could be seen as three voice calls.

6.1.2 General Model for RCST-Centric Channels

This model corresponds to RCST-centric (and not beam-centric) connectivity channels. Such channels can be defined in small mesh transparent networks or in networks supporting CC return links. They are exclusively used by an RCST, as suggested by the model in figure 6.2, which is a reflection of the general model in figure 6.1 applied to only one RCST.

An immediate consequence of limiting the channel usage to one RCST is that the bandwidth in all channels (mesh and star) originating at the RCST to different destination downlinks will never exceed the equivalent of an MF-TDMA carrier rate.

The arrangement in figure 6.2 applies to any source RCST in any UL beam (RCST A in UL Beam A is illustrated). The MPEG profile has been again considered for illustration purposes.

NOTE: Each source RCST will have its own set of independent channels.

Within each connectivity channel to a destination downlink (e.g. a TDM carrier in a DL beam), an RCST can establish multiple connections, e.g. for different C2P classes of service, to one or multiple RCSTs receiving the TDM carrier in that beam. As an example, within the connectivity channel associated with destination downlink 1, the RCST A establishes four mesh connections: two connections with C2P CoS 1 to RCST 1 and RCST 2, one connection with C2P CoS 2 to RCST 2 and one connection with C2P CoS 3 to RCST 3. It can also establish star connections with the NCC/GW in the connectivity channel to destination downlink 3, for both access traffic and M&C signalling. For destination downlink 1 the model illustrates the aggregation of connections with the same C2P/MAC CoS 1, while for destination downlink 3 the model illustrates the multiplexing of traffic connections with different C2P/MAC CoS (CoS 2 and CoS 3) under a unique PID (PID A3m).

The identification of RC/MAC classes of service and of the connectivity channels is as in the general model for beam-centric channels (figure 6.1).

The model in figure 6.2 is considered general, since it can be used for all network reference scenarios - transparent or regenerative, mesh or star. Its key feature is that the channels are RCST-specific; in this regard it is a particularization of the beam-centric general model.

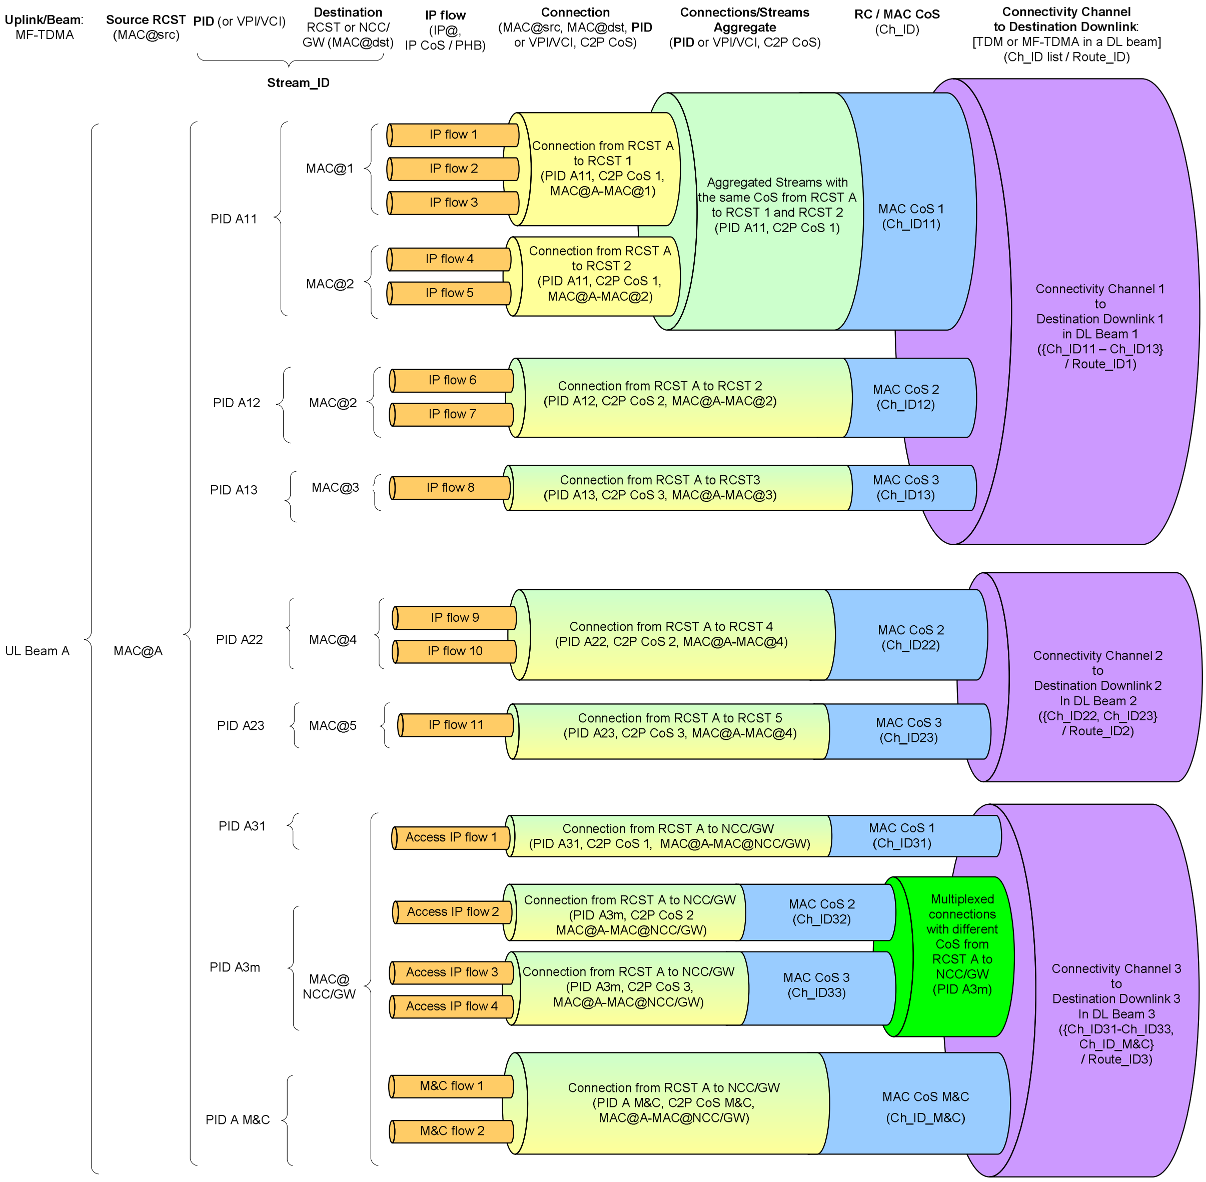

Figure 6.2: General Channel/Connection/Flow Model: RCST-Centric Channels

6.1.3 Cardinality Between C2P Elements

The cardinality diagram in figure 6.3 describes the cardinality between the C2P-related entities defined in clause 5 and captured synthetically in the general C2P models in figure 6.1 and figure 6.2.

The general C2P models are fully consistent with the SatLabs QoS models [i.4] and BSM QoS functional architecture [i.3]. It is noted that the connectivity channels are not part of these QoS models, defined for access networks, and neither are the connections. These models bring instead to the foreground the IP CoSs/PHBs and the associated support at MAC layer. The IP CoSs/PHBs are network layer functional entities vital to QoS provisioning but less important to C2P, therefore they have not been represented as separate elements in the C2P models, which are more concerned with the specific link layer elements (namely channels and connections/streams), associated with connection control and dynamic resource control. Nevertheless the PHBs can be associated with connections and RCs/MAC CoSs, as indicated in the cardinality diagram in figure 6.3. It is recalled that in the C2P model the IP CoS/PHB is seen as an IP flow attribute, which can be used to map the flow to a connection with appropriate RC/MAC CoS (as in the SatLabs QoS model/architecture).

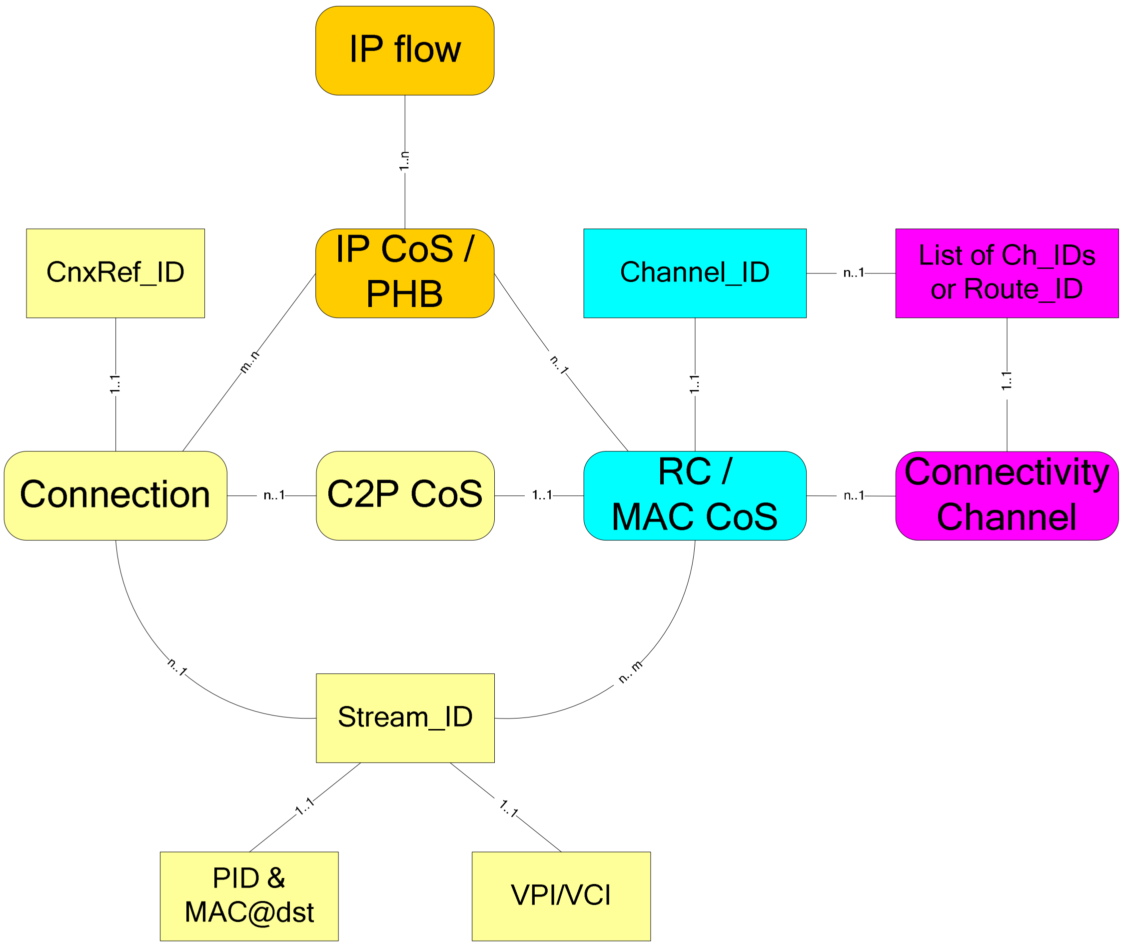

Figure 6.3: Cardinality Diagram for C2P Related Entities

Two types of objects are represented in the cardinality diagram:

- Logical entities associated with C2P (IP flow, IP CoS/PHB, Connection/Stream, RC/MAC CoS, C2P CoS, Connectivity Channel), represented as rounded boxes.
- Logical identifiers or attributes (Channel_ID, Stream_ID, Route_ID, CnxRef_ID, PID and MAC@dst or VPI/VCI), represented as squared boxes.

It is noted that C2P CoS has originally been introduced as an attribute of a connection (clause 5.4.1). Its applicability has afterwards been extended to connections/streams aggregates, therefore it is represented as a logical entity in the cardinality diagram.

The cardinality links should be understood as follows:

Object A n…m Object B

means that:

- n objects A relate to m objects B (read from A to B);
- n can take values from 1 to n;
- m can take values from 1 to m.

The cardinality between the C2P-related logical entities can be summarized as follows:

1) n…1 IP flows to IP CoS/PHB: this is in line with the IP flows aggregation (possibly with priorities) specific to DiffServ and also applicable to IntServ IP flows (see clause 5.4.1.2).

2) n…m IP CoS/PHB to Connection: n IP CoSs/PHBs (and the corresponding IP flows/IP flow aggregates) can share one connection, associated with one stream (as defined in the context of C2P); one IP CoS/PHB can be mapped to m connections (e.g. to different destination RCSTs/NCC/GW), but it is expected that all connections will share the same C2P CoS (i.e. the value of n in the n…1 connection - C2P CoS cardinality should be set to 1).

3) n…1 Connection to C2P CoS: in general each connection is defined with a C2P CoS, but this does not preclude defining multiple connections with the same C2P CoS, e.g. to different destination RCSTs/NCC/GW in the same or different connectivity channels.

4) 1…1 C2P CoS to RC/MAC CoS: each C2P CoS is mapped to one RC/MAC CoS; considering also the cardinality link 3) above, the connections with given C2P CoS in the same connectivity channels are mapped into one RC/MAC CoS (equivalent to connections/streams aggregation).

5) n…1 IP CoS/PHB to RC/MAC CoS: this is a cardinality link defined in the SatLabs QoS model; in the context of C2P, it is a consequence of the cardinality links 2), 3) and 4) above.

6) n…1 RC/MAC CoS to Connectivity channel: multiple RCs/MAC CoSs can be defined in a connectivity channel.

The cardinality between the logical entities and the corresponding logical identifiers is summarized below:

a) 1…1 Connection to CnxRef_ID: each connection is uniquely identified for C2P purposes by a CnxRef_ID and the source MAC address (if the CnxRef_ID is set by an RCST, see clause 5.4.1.1).

b) n…1 Connection to Stream_ID: several connections can share the same stream, identified for encapsulation purposes by the Stream_ID (either PID and MAC@dst or VPI/VCI). The connections can have the same C2P CoS (connections aggregation) or different C2P CoS (connections multiplexing). In the case of ATM profile the aggregation/multiplexing of connections under a unique VPI/VCI should be limited (for reasons given in clause 6.1.1) to connections having the same destination reference point - either the NCC/GW or a destination RCST.

c) 1…1 Stream_ID to PID and MAC@dst or VPI/VCI: the stream identified by the Stream_ID is as defined at the segmentation/encapsulation point; it can be an individual stream or an aggregated/multiplexed stream.

d) n…m Stream_ID to RC/MAC CoS: n streams (or streams aggregates) from different RCSTs but with the same C2P CoS, encapsulated under different PID and MAC@dst or VPI/VCI, are mapped to one RC/MAC CoS; m streams with different C2P CoS from the same RCST, mapped to different RCs/MAC CoSs, are multiplexed under the same PID and MAC@dst or VPI/VCI; m streams with different C2P CoSs from the same RCST, mapped to different RCs/MAC CoSs, may be encapsulated under different PIDs from a pool of n PID values, for the purpose of efficient capacity utilization (cross RC section packing [i.4]).

e) 1…1 RC/MAC CoS to Channel_ID; each RC/MAC CoS is uniquely identified at RCST and NCC level by a unique Channel_ID.

f) 1…1 Connectivity Channel to Channel_ID List/Route_ID: each connectivity channel is uniquely identified by a Channel_ID List/Route_ID.

As described, the cardinality diagram is consistent with the beam-centric model adopted for channels, according to which the MAC classes of service and connectivity channels are accessible to multiple RCSTs in a given uplink beam. Their identifiers (Channel_ID, Channel_ID List/Route_ID) are used in both RCST and NCC. The cardinality diagram also applies to the RCST-centric model, with the observation that the Stream_ID - RC/MAC CoS cardinality link can only be 1 : m, since all streams associated with connections with given C2P CoS (and thus given RC/MAC CoS) within a connectivity channels are assumed aggregated at the RCST into a unique stream.

6.2 Particular Models

6.2.1 Model for Mesh/Star Regenerative Networks (RSM-B Example)

The general C2P model for beam-centric channels applies to all star/mesh regenerative multi-beam scenarios, such as that exemplified by the RSM-B systems. However, since an RSM-B based system (e.g. AmerHis) was the first operational DVB-RCS system to offer mesh connectivity based on C2P, it was felt appropriate to capture and highlight in a dedicated model (represented in figure 6.4) the specificity of the RSM-B systems. The intention is to make sure that the connection control protocol specified for all DVB-RCS networks will satisfy the needs of the current and next generation RSM-B systems.

The (beam-centric) model in figure 6.4 is a customization for RSM-B systems of the general model in figure 6.1. It is recalled that an RSM-B system is based on a multibeam satellite architecture with a regenerative satellite payload, capable of supporting both mesh and star connectivity. The inter-beam connectivity is achieved by on-board table-based switching. The return links are using the MPEG profile.

As in the original model, only one uplink beam is represented (Beam A) and only two terminals (RCST A and RCST B), located in this beam. The model applies to any other RCST in any other beam.

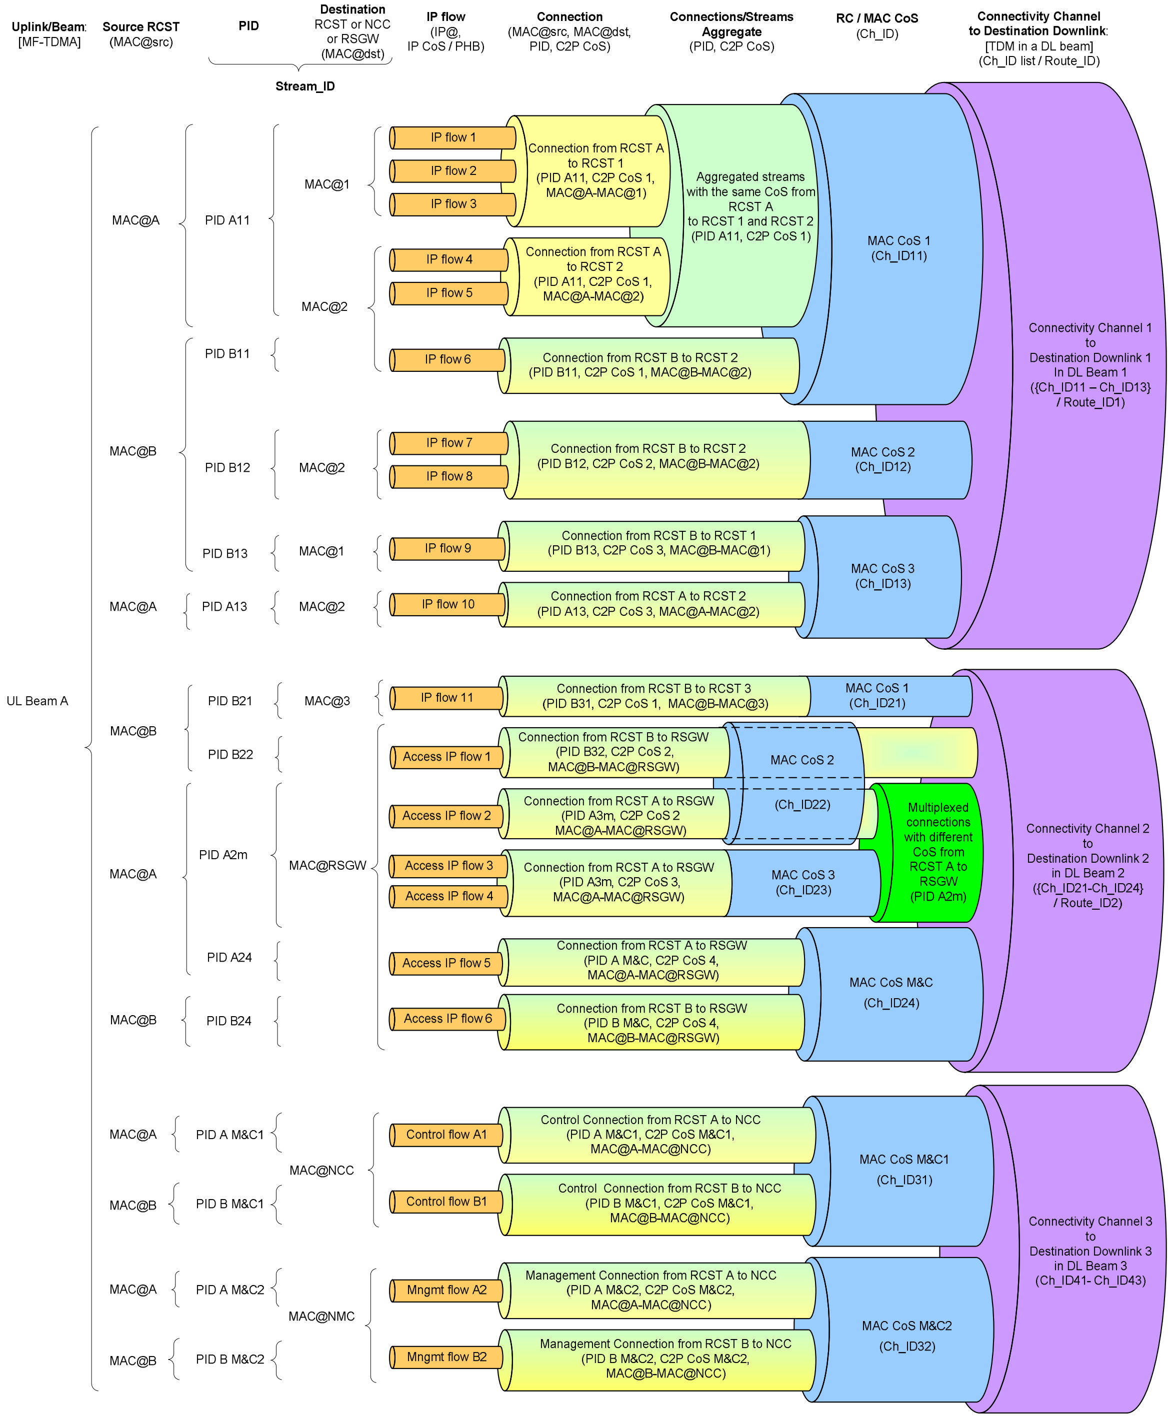

Figure 6.4. Channel/Connection/Flow Model for RSM-B Star/Mesh Regenerative Scenario

The key features of the RSM-B systems that are C2P-relveant can be summarized as follows:

- In RSM-B star/mesh regenerative systems an RCST can establish mesh connections with other RCSTs situated in the same beam or in different beams. It may also establish star connections for access traffic with the RSGW and for M&C messages with the NMC/NCC. The RSGW and NMC/NCC can be collocated or can be located in different beams, which may coincide or not with some user beams.
- Since the NMC/NCC and RSGW behave like conventional RCSTs from the point of view of the air interface and RF equipment, they transmit in DVB-RCS format and receive in DVB-S/S2 format, like any other RCST in the system. The transmission is in MPEG format, therefore the PIDs are used in the definition of connections (as connection attributes).

- The connections (both mesh and star) are established with different classes of service (C2P CoS) within connectivity channels or connection paths, associated with "fast circuits" established by the OBP, as per its switching configuration (controlled by the NMC). All connectivity channels can be regarded as mesh channels, even those that carry star connections (to RSGW and NMC/NCC), since the RSGW and NMC/NCC behave like RCSTs. Each connectivity channel is associated with a TDM, as destination downlink.

- Each C2P CoS is mapped to an RC/MAC CoS, identified by a Channel_ID. The connectivity channels will then be identified by a list of Channel_IDs/Route_ID (see clause 5.1.1).

- Within a connectivity channel the connections/streams with the same C2P CoS can be aggregated and mapped to the same RC/MAC CoS, while the connections with different C2P CoS (mapped to different RCs/MAC CoSs) can be multiplexed under the same PID.

The model illustrates various mesh and star connections established in three connectivity channels. In particular:

- In the connectivity channel 1 to destination downlink 1, the RCST A establishes 3 mesh connections, one to RCST 1 (with C2P CoS 1) and two to RCST 2 (with C2P CoS 1 and C2P CoS 2, respectively). The connections/streams with the same C2P CoS (C2P CoS1) to RCST 1 and RCST 2 are aggregated; they are sharing the same PID value (PID A11) and the same MAC CoS (MAC CoS 1, identified by Ch_ID11). The use of the same PID for two connections/streams does not raise any issue, since each destination RCST will filter the MPEG packets destined to it based on its MAC@ (inserted during encapsulation in the DSM-CC header).

- In the same connectivity channel to destination downlink 1, the RCST B also establishes 3 mesh connections: two connections to RCST 2 (with C2P CoS 1 and C2P CoS 2, respectively) and one connection to RCST 1 (with C2P CoS 3). Each connection is identified by its own PID value (PID B11, PID B12 and PID B13, respectively) and is mapped to the corresponding MAC CoS (MAC CoS 1, MAC CoS 2 and MAC CoS 3, respectively).

- The connectivity channel to destination downlink 2 is intended for star connections from the two RCSTs in the UL beam A to the RSGW (for access traffic). Five star connections are illustrated. This part of the model illustrates the multiplexing of connections with different C2P CoSs/MAC CoSs into one stream, under a common PID. In particular, the connection from RCST A to RSGW with C2P CoS 2, mapped to the MAC CoS 2 (identified by Channel_ID22) is multiplexed under the PID2m with the connection from RCST A to RSGW with C2P CoS 3, mapped to the MAC CoS 3 (identified by Channel_ID23). An additional connection to destination downlink 2 has been shown - from RCST B to RCST 3, in order to illustrate that in RSM-B systems a connectivity channel can be shared between mesh connections and star connections, in the assumption that the RSGW is located in a user beam (in which RCST 3 is located).

- The connections to destination downlink 3 are primarily for M&C signalling messages from both RCST A and RCST B to NMC/NCC. For generality, separate connections have been considered for control and for management (to NCC and NMC, respectively), with different PID values (PID A M&C1 and PID B M&C2 for RCST A; PID B M&C1and PID B M&C2 for RCST B) and different C2P CoSs (C2P CoS M&C1 for control messages and C2P CoS M&C2 for management messages). Nevertheless the M&C connections from a given RCST could be aggregated (same C2P CoS, same MAC CoS) or multiplexed (same C2P CoS, different MAC CoS), similar to the aggregation/multiplexing of traffic connections. The aggregation/multiplexing is equivalent to using only one connection for both control and management messages.

The model in figure 6.4 should be seen is an illustration of the IP flows/connections/MAC CoS/connectivity channels arrangements in RSM-B systems. Other arrangements could have been illustrated, in consistency with RSM-B specific design.

6.2.2 Model for Mesh Transparent Networks

Since a mesh transparent network is set-up as an overlay to a star transparent network, mesh channels/connections will coexist with star channels/connections (primarily used for signalling but also for access traffic). The channel/connection arrangement may be influenced by the satellite network architecture (single-beam or multi-beam).

Regardless of the satellite network architectures, it is assumed that multiple MF-TDMA carriers are provisioned in each UL beam (e.g. associated with an RL transponder).

6.2.2.1 Single Beam

All terminals are located in the same beam and are receiving the transmissions from all other terminals. The channel/connection arrangement will depend on the capabilities of the MF-TDMA demodulator in the mesh terminals.

If the MF-TDMA demodulator is a Multi-Carrier Demodulator (MCD), capable of demodulating/processing the entire RL bandwidth in the beam, the network is fully meshed (see clause 4.1.2.1) and connectivity channels are not needed. All connections (mesh and star) will be established within the unique beam. In these conditions the Channel_ID could be used with no restrictions for QoS purposes (consistent with the general model).

If the MF-TDMA demodulator is a Single-Carrier Demodulator (SCD), the network may be only partially meshed (see clause 4.1.2.1) and connectivity channels are in general needed. All RCSTs engaged in mesh communications with a given (receiving) RCST should transmit on the carrier demodulated by that RCST. From the point of view of the transmitting RCST, it should be capable of frequency-hopping among several MF-TDMA carriers, in order to establish mesh communications with multiple RCSTs (and possibly star communication with the NCC/GW). The corresponding (beam-centric) channel/connection/flow model is represented in figure 6.5, which is a particularization of the general model in figure 6.1. The particularization is with regard to the following aspects:

- The UL/DL beams are not shown any longer, since there is only one beam involved.
- ATM profile has been considered for the RL for illustration purposes, therefore VPIs/VCIs (designated as VCCs) are used instead of PIDs.
- The DL (RL) air interface is identical to the UL (RL) air interface (i.e. it is based on MF-TDMA), since the satellite payload is transparent.
- The connectivity channels are equivalent to MF-TDMA carriers.
- A separate, standard star channel (Channel_ID3) has been considered for the RL star transmissions (traffic and signalling), associated with MF-TDMA Carrier 3, but more than one carrier may be needed if the star traffic in the system is important. This should not be a problem, since the NCC/GW is typically equipped with an MCD. As an alternative, with the NCC/GW equipped with an MCD, the RL signalling and traffic could in principle be carried in any mesh channel, but this is not recommended since it would complicate DAMA scheduling and QoS management. Furthermore, the use of a standard star RL channel would be beneficial for backwards compatibility with non-mesh RCSTs and for the initial star network establishment.

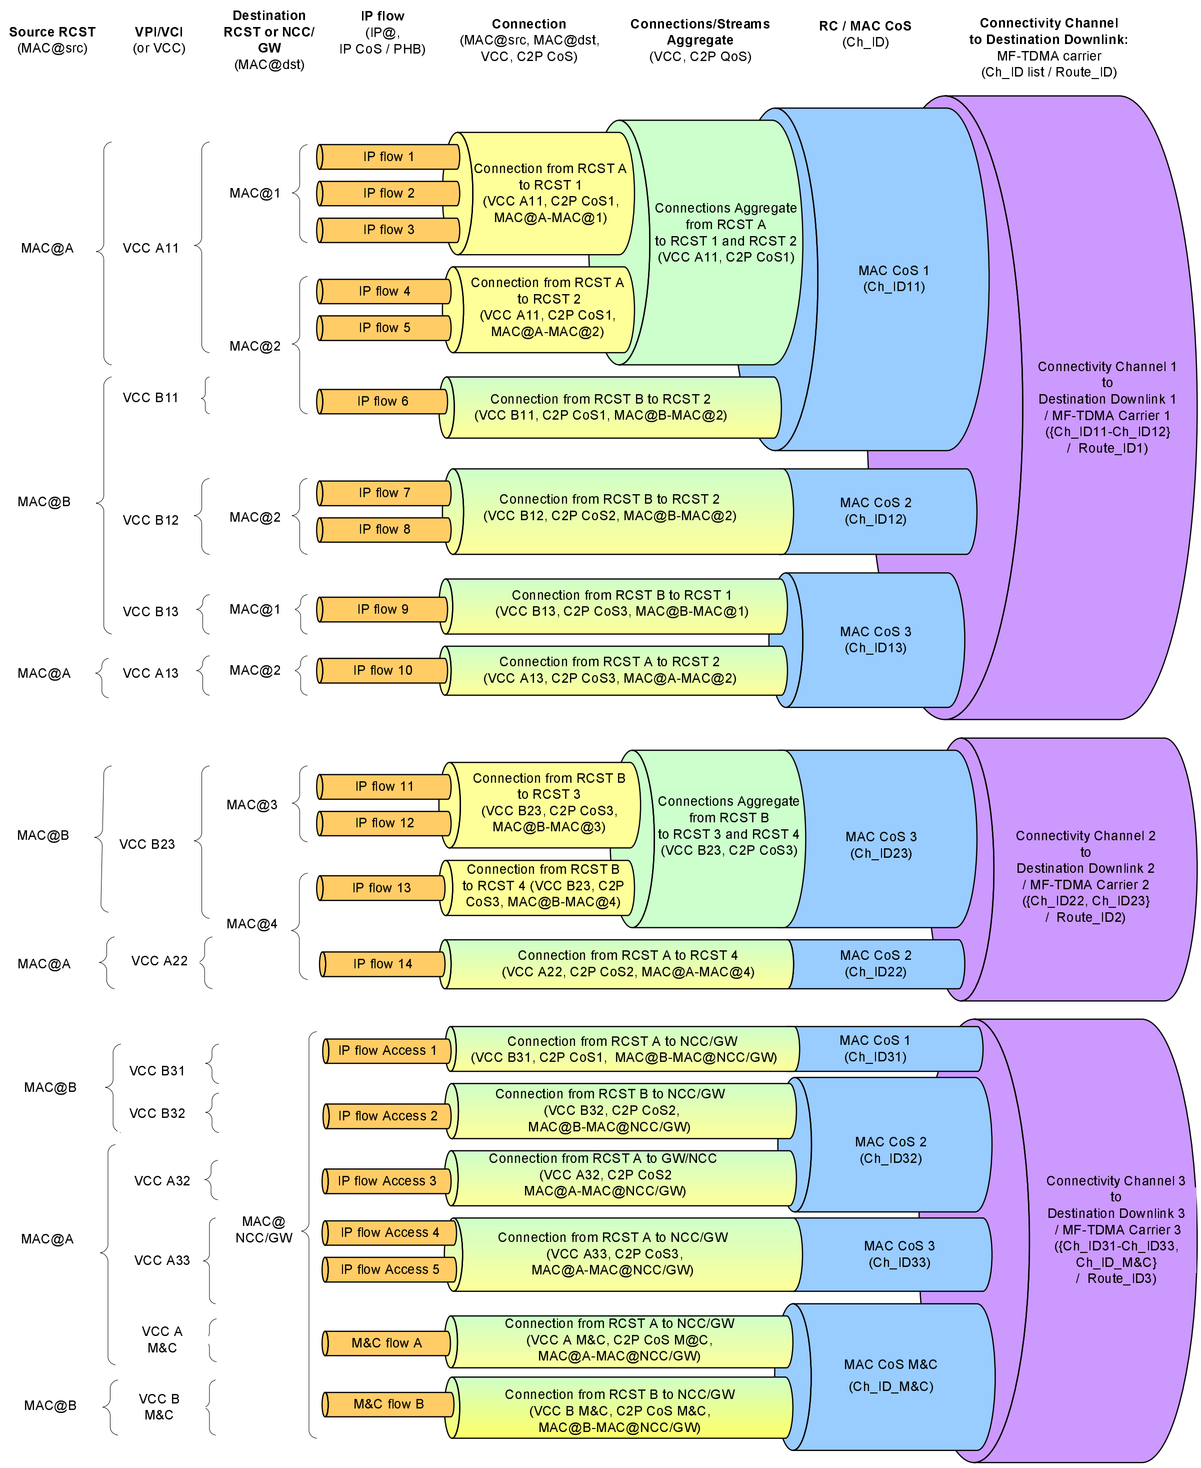

Figure 6.5: Channel/Connection/Flow Model for Mesh Transparent Scenario: Single Beam

The aggregation of connections with identical QoS requirements (same C2P CoS) to different RCSTs at the same destination downlink, has been retained in the model, but it may not be acceptable for ATM profile for security concerns, as explained in clause 5.1.1.2 (i.e. an RCST receives and can extract data intended for other RCSTs).

It is recognized that the above arrangement is not very flexible with regard to the mesh connections that can be established from a given terminal to other terminals. This is the result of having each receiving terminal "tied-up" to an MF-TDMA carrier. For more flexibility it would be desirable to have the SCD re-tuneable, but with typical existing implementations the retuning can take 2 to 3 superframes (e.g. 100 milliseconds for short superframe duration), so that in any given superframe only one carrier can be demodulated. The retuning would therefore create discontinuity in traffic receiving, unless the retuning time is accounted for in the DAMA scheduling process; this would significantly complicate the DAMA Scheduler design. An additional, non-trivial issue is that the receiving RCST should be advised on which frequency to tune in order to demodulate a burst in a given timeslot. All these problems could be circumvented and full flexibility in configuration/reconfiguration of connections could be achieved if a cheaper version of MCD were used (with a less elaborated front-end than that of the MCD used in the NCC/GW). As already indicated, in this case the connectivity channels would not be needed and the C2P model would look like that used for star transparent networks (see clause 6.2.3).

The use of MCDs or SCDs in the mesh terminal is a system/network design decision. If SCDs are used, the connectivity channels will be quasi-static, as long as the retuning cannot be achieved within a superframe. Since a mesh channel is equivalent to an MF-TDMA carrier and there are in general more terminals than carriers in the system, the channels needs to be shared among terminals; this justifies the adoption of the beam-centric channels.

Other factors to consider in mesh receiver design are the possible need for simultaneous reception from two or more mesh terminals and the need to support two or more symbol rates for fade mitigation, in the case where the transmissions from other terminals were affected by signal fading on mesh carriers.

6.2.2.2 Multiple Beam

Mesh communication between RCSTs in different beams assumes on-board connectivity between beams. In the case of a transparent satellite the connectivity is provided at the physical layer and consists in routing ("channelizing") uplink MF-TDMA carriers or group of carriers in a beam to different downlink beams. Since the mesh network is an overlay to the star network, the Gateway should provide forward links (i.e. TDMs) to each downlink beam for FL signalling and also for FL access traffic, if any.

Since the on-board connectivity is synonymous to carrier routing, the granularity (and therefore the mesh channel) is limited to one MF-TDMA carrier.

NOTE: Sub-carrier granularity will require time/frequency switching, as in the case of WEST [i.30]. system for example, which cannot be achieved by a simple transparent satellite.

As in the case of the single-beam architecture, the mesh channels will be shared by multiple RCSTs (i.e. they are beam-centric). In these conditions the general model in figure 6.1 could be used, with the observation that a mesh channel is associated with an MF-TDMA carrier or carrier group (and not a TDM) in a destination downlink beam.

6.2.3 Model for Star Transparent Networks

Assuming a single NCC/GW per satellite network, there is no need for connectivity channels in this scenario, since all RL transmissions end up in the same service beam where the NCC/GW is located. The GW should be capable of processing the RL and FL traffic from/to all beams. It therefore needs to be equipped with multiple MCDs and have at least one TDM in each beam.

An uplink beam can be outfitted with more than one transponder and an RCST is assigned to one of these transponders. Any reference to an UL beam should be understood as a reference to a transponder in that beam. The RL resources in each uplink beam/transponder are managed independently from the RL resources in any other uplink beam/transponder.

Since the Channel_IDs are not needed for connectivity in this scenario, they can be used entirely for QoS differentiation, as in the case of SatLabs reference network, which is typical of this scenario. The corresponding channel/connection/flow arrangement is represented in figure 6.6 for UL beam A, but is applicable to any other UL beam in the system. It can be seen as a particular case of the general arrangement in figure 6.1, from which the mesh channels have been removed.

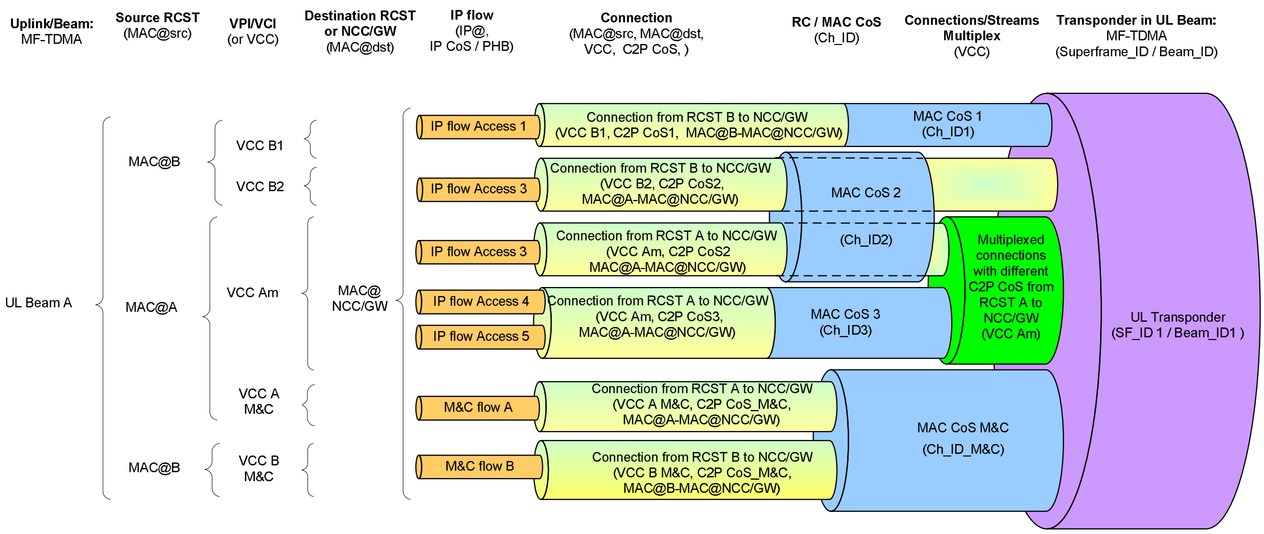

Figure 6.6: Channel/Connection/Flow Model for Star Transparent Scenario

The following comments can be made, when comparing this arrangement with the general model in figure 6.1:

- The connectivity channels are no longer needed and are replaced by the full UL transponder/beam.
- An UL beam is uniquely identified by the Beam_ID. A beam can be assigned multiple transponders. The resources in each transponder are typically associated with a Superframe_ID (SF_ID). At logon time an RCST is assigned to one of the transponders available in the beam, therefore the arrangement above is for an SF_ID within a Beam_ID.
- ATM RL profile is illustrated.
- The destination address is the same for all connections (MAC@NCC/GW).
- Since the destination is unique, there is no aggregation of connections/streams (with the same C2P CoS) within the RCST. However, the multiplexing of connections/streams with different MAC CoS can be supported, regardless of the RL profile (ATM or MPEG).

7 C2P Network Interworking Aspects

7.1 C2P Triggers

This clause identifies various events that can trigger C2P messages for connection establishment, modification or release.

The list below includes a number of events/C2P triggers identified so far. Other triggers may be considered in the future.

1. Arrival of an IP packet at the RCST (leading to a connection establishment/modify request). Three triggers can occur, following the comparison of the information in the packet header with a filter criteria/mask:
   - addressing/routing trigger (for connection establishment), if the packet matches an existing flow type (with defined IP CoS/PHB), but its IP destination address does not match any of the current connections;
   - QoS trigger (for connection establishment), if the packet's IP CoS/PHB does not match any existing flow type;
   - flow trigger (for connection modification), if the packet matches an existing flow type (with defined IP CoS/PHB) and the IP destination address of an existing connection, but the packet is not part of a flow carried by the connection.
2. Expiration of a connection inactivity period (leading to a connection modify/release request).

3. Traffic measurement at RCST or Hub (leading to a connection modify request).
   - On RCST side the trigger can be based on the capacity occupancy of a buffer (queue size) or on the number of IP flows per connection.
4. Changes in the network loading (leading to a connection modify request).
   - It may happen that when a connection was established at RCST's request, the NCC could not assign the entire capacity requested (within subscriber's SLA profile). When capacity becomes available, the NCC may initiate a C2P message to grant additional capacity (up to the subscriber's SLA).
5. Changes to subscriber's SLA, as a result of SLA renegotiation (leading to RC modify request and possibly to connection modify requests).
6. Signalling messages on the interface with other network mechanisms, e.g. Performance Enhancement Proxy (leading to a connection modify request).
7. Interception of a signalling message associated with an application session (e.g. RSVP, SIP) (leading to a connection establishment/modify/release request).

   This type of trigger assumes a control plane interface between the application layer and the transport layer.

The above events can be further analysed in order to differentiate between RCST triggers and NCC triggers.

The events listed under items 6 and 7 will not be considered in the present document and in the first revision of C2P specification.

NOTE: Some of the above triggers, associated with subscriber's SLA and possibly with C2P CoS, may have as effect the modification of the parameters associated with a Request Class. This should not be considered a channel modification command. A true channel modification command refers to changing in the DAMA Scheduler the parameters of a bandwidth partition associated with a MAC CoS. Such a channel modify command cannot be explicitly sent by RCSTs; it could instead be local to NCC, which is responsible for the partition of return bandwidth among MAC CoSs. The command can be triggered by C2P requests from RCSTs, e.g. when the capacity requested for the connections mapped to an RC/MAC CoS approaches the boundaries set for that MAC CoS, while the partitions for other MAC CoSs are underutilized. It can also be triggered by explicit requests for changes to subscriber SLA. The NCC policies for channel modification (i.e. for changing the return link partitions) are system specific and are out of scope of this report.

7.2 C2P/QoS/Resource Control Interactions

The purpose of this clause is to describe the interactions between the QoS models, the MAC layer support for QoS provisioning (resource control/DAMA scheduling) and the connection control functionality/entities. More specifically, the clause describes the tables mechanism defined for the linkage/associations between various parameters that need to be configured at IP layer and MAC layer for traffic plane processing, in the context of the dynamic connectivity provided by C2P, which performs a control plane function. This is preceded by a review of the QoS model adopted for the purpose of C2P definition/specification.

7.2.1 C2P QoS Models

7.2.1.1 Return Link C2P QoS Model

As stated in clause 5.4.1.2, the MAC layer connections, as defined in the context of C2P, are established with predefined QoS/classes of service, referred to as C2P CoSs, which actually correspond to the MAC CoSs implemented in the DVB-RCS networks, as part of return link resource control. This implies that the QoS model considered in this clause are applicable to the return link. Furthermore, since the RCST is the return link ingress point into the DVB-RCS service domain, the QoS model actually applies to the RCST.

For the approach to QoS provisioning on the forward link please refer to clause 7.2.1.2.

The IP packets entering the RCST, before being mapped to connections with predefined C2P CoSs, are first associated with various IP classes of service (IP CoS).

The QoS model adopted for C2P definition/specification supports both the DiffServ (RFC 2474 [i.14] and RFC 2475 [i.15]) and IntServ (RFC 1633 [i.16]) architectures (frameworks or models). The C2P QoS model is therefore consistent with the SatLabs QoS model [i.4], based on the DiffServ architecture, and also with the BSM QoS model [i.3], supporting both the DiffServ and IntServ architectures. The IP queuing structure from the BSM QoS model, which is also adopted for the C2P QoS model, is reproduced for convenience in figure 7.1.

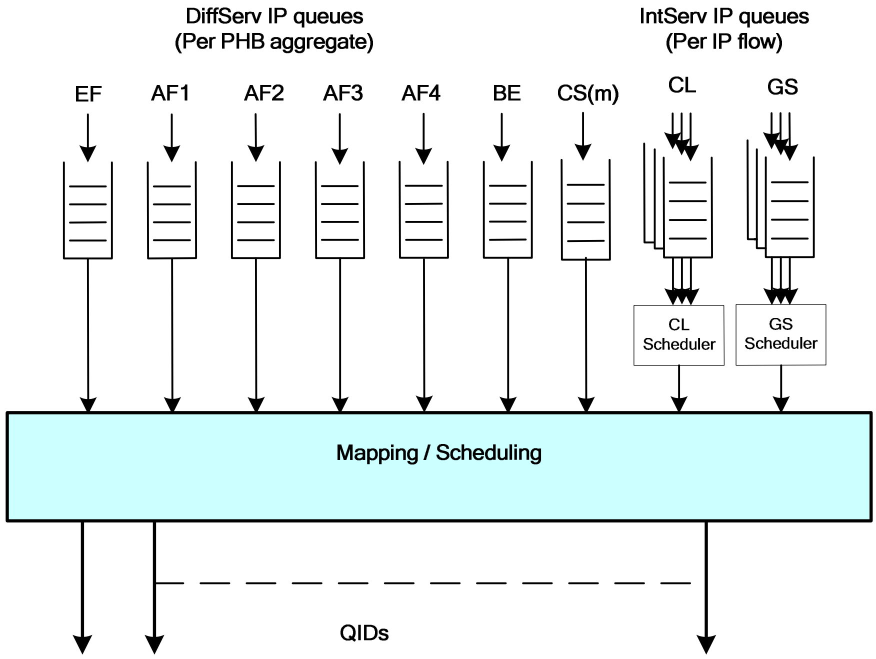

Figure 7.1. RCST IP Queue Architecture

Figure 7.1 captures the key features of the two frameworks, namely:

- In IntServ separate IP queues need to be maintained for each IP flow, associated with either the Guaranteed Service (GS) class (RFC 2212 [i.17]) or Controlled Load (CL) class (RFC 2212 [i.17]).
- In DiffServ separate IP queues are only required for IP flow aggregates, associated with the DiffServ PHBs. All standard PHBs (EF, AF1-4, BE) and a generic Class Selector PHB (CS(m)) are illustrated. Nevertheless separate queues can optionally be maintained for individual IP flows (as in IntServ) or groups of IP flows within a PHB aggregate, e.g. in order to provide priorities or fairness among flows within the PHB aggregate.

Figure 7.1 also illustrates the mapping (via scheduling processes) of the IntServ IP flows into two abstract queues at the SI-SAP interface, identified by their QIDs. Similar schedulers would be needed if per-flow queues were considered for the DiffServ architecture.

NOTE: One can even conceive scheduling of both IntServ and DiffServ IP flows into a unique QID, but this is difficult to manage and therefore not recommended.

The purpose of mapping/scheduling of the IP flows into a reduced number of abstract queues is to limit the number of MAC Classes of Service that need to be supported in the DVB-RCS network.

A more detailed RCST functional QoS architecture is represented in figure 7.2.

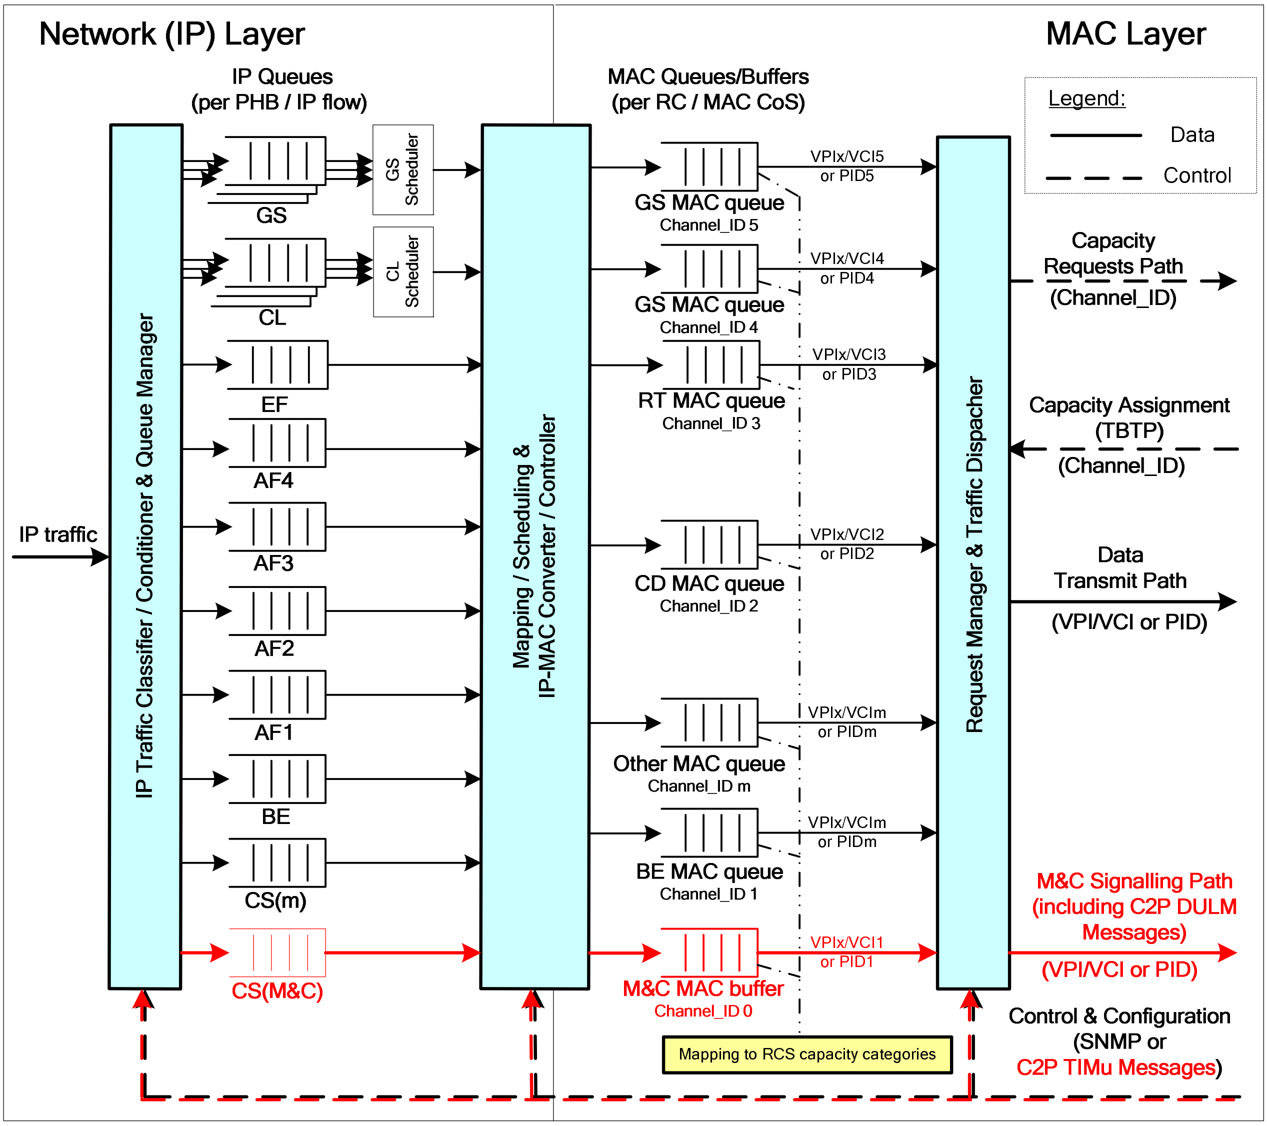

Figure 7.2. RCST QoS Functional Architecture

The architecture includes two sets of queues, one at IP layer and another one at MAC layer, and a number of other functional blocks.

At IP layer the following queues are illustrated, for generality:

- One queue for each standard DiffServ PHB (EF, AF1-4, BE), but not for the individual flows in the aggregates, since this is not a DiffServ requirement.
- One queue for a generic Class Selector PHB, identified as CS(m).
- A separate queue, labelled as CS(M&C), for all management and control (M&C) signalling messages, generically designated as M&C messages. Such messages can be external to the DVB-RCS network, e.g. associated with Inter-Network Management (INM) or with protocol signalling (e.g. application session protocols, QoS signalling protocols), or internal to the DVB-RCS network, for Local Network Management (LNM), e.g. OAM messages, C2P messages, etc. The C2P messages are carried in Information Elements (IEs) in TRF bursts, based on the DULM method [i.1].
- Two sets of per-flow queues for IntServ classes of service (GS and CL). The IP flows in each class are aggregated based on a defined scheduling discipline (e.g. weighted fair queuing) before being sent to the MAC layer via SI-SAP interface.

The forwarding rules for packets in all IP classes of service but the CS PHBs are defined by the IntServ and DiffServ standards. For CS PHBs the forwarding rules are not standardized and are in general system/network specific. As an example, the forwarding rules for the M&C messages are most probably defined by the network operator. Nevertheless all CS PHBs used in a system should be defined to be consistent with the DiffServ specifications (RFC 2774 [i.31]).

A reduced set of queues are illustrated **at MAC layer**, namely:

- one queue for each of the IntServ class of service (GS and CL);
- three queues for the three RCs recommended by SatLabs for short-term deployment, i.e. for Real Time (RT), Critical Data (CD) and Best Effort (BE) traffic;
- one queue, labelled as "Other MAC Queue", added for generality;
- one queue for M&C signalling.

The MAC queues are also referred to as buffers, since the queuing is expected to take place primarily at the IP layer. Moreover, depending on the implementation, the IP to MAC conversion can be achieved "on the fly", in which case the MAC queues/buffers would not be needed. However, they have been represented in figure 7.2 as functional blocks and could be seen as abstract queues, identified for the purpose of DAMA scheduling by Channel_IDs (corresponding to the QIDs in the BSM QoS model). The MAC layer packets from each queue are typically associated with a separate stream and encapsulated under the same VPI/VCI or PID. However, for consistency with the SatLabs QoS model and C2P models in clause 6, the RCST QoS model does not preclude the multiplexing of packets from different RCs/MAC CoSs into a unique stream and their encapsulation under the same VPI/VCI or PID. This is illustrated in figure 7.2 by the use of VPIx/VCIm (or PIDm) for the encapsulation of packets from both the BE MAC queue and Other MAC queue. The other VPI/VCI or PID values have been arbitrarily labelled, and so have the Channel_ID values, with the exception of Channel_ID 0 which is assigned by default to the signalling channel. However, the assumption was made that all virtual circuits (identified by different VCI values) are within the same virtual path, identified by VPIx.

The **IP Traffic Classifier/Conditioner & Queue Manager** is in charge of assigning the incoming IP packets to one of the IP queues and conditioning them accordingly. More particularly, it is responsible for:

- multifield Traffic Classification, based on filter patterns/masks;
- steering the IP packets to the appropriate queues;
- traffic Conditioning (metering, policing, shaping, dropping);
- managing of IP queues, including setting the drop precedence and handling packet dropping.

Traffic conditioning and queue management are performed according to a set of rules established for each IP queue, consistent with and as an enforcement of the Traffic Conditioning Agreement (TCA) established between a subscriber and its Service Provider.

The **Mapping/Scheduling & IP-MAC Converter/Controller** function performs two related tasks at the SI-SAP interface:

- mapping of packets from different IP queues into a limited number of MAC CoSs/queues;
- conversion of IP packets into the MAC layer format specific to the satellite return link (ATM or MPEG).

The mapping function relies on one or several scheduling algorithms, depending on the number of CoSs/PHBs mapped into one MAC CoS. The mapping is flexible, in full consistency with the cardinality model described in clause 6.1.3. A good practice is to use 1:1 mapping for IP packets associated with performance-sensitive IP CoSs/PHBs. As illustrated in figure 7.2, one MAC CoS is suggested for GS, CL and EF traffic and for M&C signalling messages, respectively, while all AF packets can be mapped into one MAC CoS/queue, e.g. the CD MAC queue. Likewise, the BE MAC queue can be used for IP packets from the BE PHB and other PHBs. In general, the IP packets from IP CoSs/PHBs with similar service requirements are mapped to a unique MAC CoS/queue. This will limit the number of MAC CoSs that need to be implemented in the DVB-RCS network.

The **Request Manager & Traffic Dispatcher** handles two related functions associated with the DAMA operation:

- calculation and transmission of capacity requests. Separate requests are issued per MAC CoS, for each DVB-RCS capacity category (i.e. RBDC or VBDC) authorized for the MAC CoS.

- scheduling/dispatching of packets from different MAC queues in the slots assigned via the TBTP.

All the functions described above are transport stratum user-plane functions, with the exception of capacity request calculation/transmission, which is a control plane function, part of DAMA scheduling. The RCST behaves as a DAMA client, while the DAMA server (i.e. DAMA Scheduler), responsible for the dynamic capacity assignment, is located at the NCC.

Figure 7.2 also illustrates other management and control plane functions, namely the configuration of the QoS mechanisms at IP layer and MAC layer with adequate QoS parameters and polices. These functions are part of the application stratum in the overall multi-strata QoS management architecture [i.3]. The configuration can be static or dynamic. Dynamic configuration can be achieved by:

- using the harmonized management plane tools recommended by SatLabs (based on SNMP and QoS-related MIB objects);

- relying on the C2P forward link signalling in the case of systems supporting dynamic connectivity. C2P messages are carried in Information Elements in the unicast TIM (TIMu).

Configuration parameters and policies are required for each IP layer queue and each MAC layer CoS/queue. The policies are in general captured in the form of algorithms, which include a number of fixed parameters and a number of variable (configurable) parameters. The policies, and therefore the algorithms, are class-specific. They are also system/network specific. However, for the sake of interoperability they should rely on a number of common, configurable parameters. In a system supporting dynamic connectivity such parameters can be transported in C2P messages. At the RCST level the parameters are included in a set of interlinked data structures (tables), defined for configuration at various levels (clause 7.2.3).

C2P itself is part of the control plane functionality in the application stratum of the overall QoS model. Its main responsibility is to establish connections with defined quality of service (i.e. C2P/MAC CoS), and to modify/release them. The connections themselves are not shown in figure 7.2, which only represents the transport stratum functional blocks in the overall QoS model. Providing the resources required to meet connections' QoS requirements is achieved by proper reflecting the QoS connection profile (consistent with its C2P CoS) into the configuration parameters for the transport mechanisms (at both IP and MAC layer). It is only this configuration aspect of C2P functionality that is captured in the RCST functional diagram in figure 7.2. The interdependencies between C2P entities and QoS transport stratum entities are as per the C2P channel/connection/flow models and cardinality model in clause 6.

7.2.1.2 Approach to Forward Link QoS Provisioning

The approach to QoS provisioning on the forward links from NCC/GW to RCSTs is also based on the DiffServ PHBs (and/or IntServ CoS), but the QoS model is rather different from the return link model, since the forward link bandwidth resources are not dynamically managed (i.e. based on DAMA) but rather time shared among satellite terminals. Furthermore, the QoS model applies to a much higher level of traffic and connection aggregation taking place in the NCC/GW. The IP CoSs/PHBs are mapped to logical partitions of the TDM bandwidth. The partitions are defined at different levels (groups of RCSTs, RCSTs, IP subnets) and configured with min/max rates and various policies.

7.2.2 Required Configuration Parameters at IP and MAC Layers

7.2.2.1 General

At IP layer the configuration parameters are used to configure the DiffServ and IntServ mechanisms used for IP packet classification/conditioning and for packet queuing/scheduling. The configuration parameters are class specific, since the behavioural/forwarding requirements are class specific. Furthermore, the implemented mechanisms used to satisfy the forwarding requirements are in general manufacturer specific.

One could argue that IP QoS provisioning is part of the Satellite Independent Access Function (SIAF), therefore it should not be part of any further regulation. However, a minimum set of configuration parameters needs to be identified to support the basic forwarding rules defined by standards for each class, in order to establish a basis for interoperability.

7.2.2.2 IP Forwarding/Behavioural Rules

In the context of the **DiffServ** standard, the key behavioural features of the IP aggregates/PHBs are as follows:

- The **EF class** (RFC 3246 [i.34]) provides guaranteed rate for in-profile traffic and in this respect it is similar to a "virtual leased line". It is used for end-to-end services with low loss, low latency, low-jitter and guaranteed bandwidth. Such services, also described as premium services, are intended for high priority real-time traffic.

- The **AF classes** (RFC 2597 [i.33]) are used for end-to-end services that need assurances of high probability of delivery of packets within a subscription profile (booked bandwidth). The four defined AF classes are differentiated by the amount of allocated forwarding resources (bandwidth and buffer space), but there is no defined priority/precedence among classes. The level of forwarding assurance will thus depend on the allocated resources, the current load and, in case of congestion, on the drop precedence.

The above forwarding specification suggests a guaranteed rate for the traffic within the subscription profile, and a best effort service for the excess traffic.

- The **BE class** is used for end-to-end services that have no performance and bandwidth requirements. BE traffic aggregates may be subject to flow control or dropping policies.

The above description is complemented by the requirement that the traffic in one class should not negatively impact the traffic forwarding in any other class. The following precedence among PHBs should be observed: EF (highest), followed by AF and then by BE (lowest).

In the context of the **IntServ** standard, the key behavioural features for the IP flows are as follows:

- The **GS class** (RFC 2212 [i.17]) provides guarantees that the in-profile traffic will arrive within the guaranteed delivery time and will not be discarded due to queue overflow. There is no guarantee on jitter.
  GS class corresponds to virtual circuits (established using an end-to-end signalling protocol and admission control) and is used for applications requiring low latency and guaranteed delivery (e.g. audio/video play-back applications).

- The **CL class** (RFC 2211 [i.32]) provides a QoS closely approximating the QoS that the same flow would receive from a lightly loaded network, but uses admission control to ensure that the service is provided at the desired rate even when the network is loaded. There are no specific performance requirements and it is acceptable to have a small percentage of packets undelivered or delayed by much more than the delay experienced in the unloaded network.

7.2.2.3 C2P-Related Configuration Parameters

Based on the behaviours described in 7.2.2.2, the following parameters can be specified for the DiffServ PHB aggregates:

For EF class:

- Guaranteed rate, typically defined as Peak Data Rate (PDR).

- Buffer size, which together with the guaranteed rate can fully define the in-profile traffic (optional, depending on the implementation of the conditioning mechanism, e.g. as single-rate token bucket).

- Delay specification (may become relevant for system with long frame duration; it is system dependent).

- Jitter specification (relevant for systems with long superframe duration).

NOTE 1: If guaranteed static capacity is used for PDR and the superframe is short (e.g. < 50 milliseconds), the delay and jitter are of no concern.

For AF classes:

- Guaranteed rate, for the subscription (booked) rate. This can be defined as Sustainable Data Rate (SDR).

- Non-guaranteed rate, defined slightly above the guaranteed rate but below the maximum transmit rate. It is used, together with the guaranteed rate, to establish the drop precedence levels.

NOTE 2: If SDR is used for the guaranteed rate (which is more common), the PDR parameter may be used for the non-guaranteed rate. This will prevent the introduction of additional new parameters, observing that the meaning (and usage) of SDR and PDR is class specific. In the case of AF class for example, (PDR-SDR) actually corresponds to the non-guaranteed portion of the traffic.

- Buffer size, which together with the two rates can fully define the levels of packet non-conformance (optional, depending on the implementation of the conditioning mechanism, e.g. as two-rate token bucket).

For BE classes:

- Non-guaranteed rate, which can be defined as Average Data Rate (ADR) or Minimum Data Rate (MDR).

NOTE 3: ADR or MDR are optional. The already defined SDR can be used instead, but with a BE-specific meaning.

Based on the behaviours described in clause 7.2.2.2, the following parameters can be specified for individual IP flows associated with IntServ classes of service:

For GS class:

- Guaranteed rate, defined as Peak Data Rate (PDR) (with the same meaning as for EF class).
- Buffer size, which together with the guaranteed rate can fully defined the in-profile traffic (optional, depending on the implementation of the conditioning mechanism, e.g. as single-rate token bucket).
- Delay specification (optional).

NOTE 4: A minimum delay could be imposed to make sure that the end-to-end delay, committed during path reservation, is guaranteed.

For CL class:

- Guaranteed rate, e.g. defined as Sustainable Data Rate (SDR) or Peak Data Rate (PDR).

NOTE 5: The rate is expected to be soft-guaranteed by the network for the estimated rate provided by the client, i.e. the network may allow statistical multiplexing. For the traffic within the negotiated rate the client can except fairly constant delay and (low) drop ratio.

The above IntServ behaviours and parameters are defined for individual IP flows. If aggregation of IP flows in a class is used, as indicated in the QoS model for both GS class and CL class, the parameters need to be aggregated at class level. If statistical multiplexing is allowed, the aggregate rate could be smaller than the sum of individual flow rates. A scheduling discipline is therefore needed to provide fairness and/or priorities among flows. As an example to illustrate the effect of statistical multiplexing, if 20 VoIP flows are aggregated, the required bandwidth would be only about 50 % of the bandwidth which would otherwise be required if no aggregation were used.

As a result of flow aggregation, only the IntServ classes of service (and not the individual flows in each class) will be mapped to MAC CoSs, similar to the mapping of the DiffServ PHBs.

Priority can also be provided among the IP flows or groups of IP flows in the DiffServ PHB aggregates. In this case, separate queues need to be maintained/configured per IP flow or group of IP flows in a PHB (for the purpose of traffic conditioning), similar to the IntServ case; the flows will then be aggregated (based on a scheduling discipline) at PHB level, while allowing for statistical multiplexing gain. This is not part of the baseline and in anyways will only have impact on IP QoS functionality and not on the connection definition/handling at the MAC layer.

The above configuration parameters are part of the IP flow profiles (for IntServ) or of the PHB profiles (for DiffServ) and are primarily used for packet conditioning. They are associated with the bandwidth (or rate) parameters of the connections (consistent with their C2P CoS), but the association is not straightforward, as it will be shown in clause 7.2.3; it has to be consistent with the C2P model and the cardinality model, which are rather complex. The connections' rate parameters will then be mapped to capacity limit values of the corresponding RCs/MAC CoSs. The mapping is needed for DAMA operation and is performed at the NCC/GW.

For the connection associated with the IntServ IP flows, the connection profile should only reflect the aggregate rate parameters at the IP CoS level (and not at IP flow level), which is used for IP flow scheduling but not for any further conditioning of IP packets.

C2P should not be concerned with the configuration of parameters at IP flow level. This can be achieved for all defined flow types by a management function, e.g. by configuration at RCST's initialization or via SNMP messages. C2P should only be concerned with the dynamic configuration of rate parameters at IP CoS/PHB level, reflecting the number of active flows of given type in the IP aggregate, and of the capacity limit values at the RC/MAC CoS level. The dynamic configuration is the result of C2P request/response messages for connection establishment or modification.

The relationship between the above C2P-related parameters are summarized in table 7.1 and are IP CoS/PHB specific. The assumption is that only two rate parameters are used as part of the IP flow profile - the SDR and PDR, whose meaning depends on the class of service. The table also includes recommendations/guidelines for the mapping of rates to RCS capacity limit values for all defined IntServ and DiffServ classes of service.

Table 7.1: C2P-Related Bandwidth Parameters

IP CoS/PHB	Per IP flow request rate parameters (Flow profile)	Per IP CoS/PHB request/response rate parameters	Per RC response capacity parameters	Comments
IntServ GS class	PDR_{flow} $SDR_{flow} = PDR_{flow}$ Buffer size Delay (optional)	PDR = $\alpha\sum PDR_{flow}/PDR'$ SDR = PDR Buffer size	CRA = PDR' RBDCMax = 0 VBDCMax = 0	PDR' is understood as true peak rate (in-profile traffic), which is guaranteed α: flow multiplexing gain ($0.5 \leq \alpha \leq 1$)
IntServ CL class	SDR_{flow} $PDR_{flow} = SDR_{flow}$	SDR = $\beta\sum SDR_{flow}/SDR'$ PDR = SDR	CRA = 0 RBDCMax = SDR' VBDCMax = 0	SDR' is understood as guaranteed rate; could also be mapped to CRA β: flow multiplexing gain ($0.3 \leq \beta \leq 1$)
EF PHB	NA	PDR/SDR' SDR = PDR Buffer size	CRA = $\sum$PDR' RBDCMax = 0 VBDCMax = 0	PDR' understood as true peak rate (in-profile traffic), which is guaranteed
AFn PHB	NA	PDR/PDR' SDR/SDR' Buffer size	CRA = 0 RBDCMax = $\gamma\sum$SDR' VBDCMax = $\gamma\sum$(PDR'-SDR')	SDR' is understood as guarantee (booked) rate; could also be mapped to CRA PDR' is understood as non-guaranteed rate: SDR' < PDR' < Max RCST rate γ reflects n:1 PHB to RC mapping; can be < 1 if statistical multiplexing gain is sought
BE PHB	NA	SDR/SDR' PDR = SDR	CRA = 0 RBDCMax = 0 VBDCMax = $\delta\sum$SDR'	SDR is understood as non-guaranteed rate (averaged per superframe) δ reflects n:1 PHB to RC mapping; can be < 1 if statistical multiplexing gain is sought

NOTE 1: All parameters apply to the uplink return path, therefore they are transmit parameters. In the case of bidirectional connection rate parameters will also be defined for the downlink of the forward path (receive parameters); they will only be converted in capacity limit values for mesh bidirectional connections.
NOTE 2: SDR/PDR are per aggregate rates requested by RCST.
NOTE 3: SDR'/PDR' are per aggregate response rates, following the execution of admission control in NCC: PDR' ≤ PDR, SDR' ≤SDR'.
NOTE 4: With the exception of AFn PHBs, for all other PHBs a 1…1 PHB to RC mapping was considered. In the case of AF, multiple PHBs (n = 1-4) are typically mapped to one RC, therefore the RC capacity limit values should reflect the aggregation (sum) of individual rates, possibly with an allowance for statistical multiplexing ($\gamma \leq 1$).
NOTE 5: The sum of all capacity limit values for all request classes should not exceed the RCST transmission rate.
NOTE 6: One RC per IP CoS/PHB assumed for illustration purposes, but in general several IP CoSs/PHBs could be mapped to one RC.

Even if not shown in the RCST QoS functional model, the connection requests can be triggered by the detection of a new IP flow or a new IP flows aggregate (among other C2P triggers, see clause 7.1). The requests can be for a new connection or for the modification of an existing connection. Connections with the same C2P CoS are customarily mapped to the same RC.

7.2.3 RCST Data Structures for Dynamic Connectivity

7.2.3.1 Overview

Dynamic connectivity is used primarily to setup connections with defined C2P CoS/QoS profile and to modify/release them. Within this broad scope C2P also performs address resolution for the purpose of MPE encapsulation (i.e. the association between the MAC addresses and IPv4 addresses of the parties involved in a connection). Setting-up/modifying connections involves dynamic bandwidth allocation (following the execution of an admission control function) and configuration of parameters in various network components. With regard to this last aspect, C2P can be used for the dynamic configuration of RCST parameters based on a number of RCST data structures. Some parameters in the data structures are not used in the C2P messages but only for local (internal) processing.

In order to support the full functionality assumed by C2P, the RCST contains the following internal data structures: the Packet Classification Table, PHB Mapping Table and Request Class Table, already defined in the SatLabs SSR - Part 3 [i.7] (subject to modification, to reflect the more complex C2P QoS model and other C2P needs), and two new tables - the Active Connection Table and the Connection Profile Mapping Table. Each table is briefly described below and than detailed in the following clauses.

The **Packet Classification Table** defines the classification of packets entering the RCST into packet types or flow types, based on a set of parameters defining filter masks/criteria.

The **PHB Mapping Table** defines the PHB parameters for each PHB entry and provides the association of a packet of given type, related to a classification criteria, to a PHB. Each PktClass is mapped to a PHB Entry, referenced by a PhbIndex, and each entry provides the PHB mapping to a Request Class.

The **Request Class Table** defines all layer 2 QoS parameters for each supported RC, i.e. the MAC layer logical parameters (Channel_ID, PIDs or VPI/VCI) and bandwidth parameters, expressed as capacity categories limit values (CRA, RBDCMax, VBDCMax). These parameters are controlled by the NCC and configured in both the NCC and RCST; in RCST they are used for capacity request calculation and possibly in the traffic dispatching process.

The above three tables are configured by management with initial values. The last two tables will then be dynamically updated by relying on C2P messages, every time a new connection is created/modified or released.

The **Connection Profile Mapping Table** (new) includes a set of connection QoS/bandwidth parameters, used to determine the contents of the corresponding C2P message..It is configured by management.

The **Active Connection Table** (new) contains all necessary data to fully characterize all open/active connections of a given RCST, whether initiated by the RCST or by the NCC . Each entry in the table corresponds to an active MAC connection, is identified by a Connection Reference Identifier (CnxRef ID) and contains all the parameters associated with that active connection, obtained from the connection profile parameters in the Connection Profile Mapping Table and/or from NCC via C2P response messages.

Each table will in the end be formally specified as a set of objects in tables of a C2P subgroup under the dvbrcsifMIB [i.7].

The tables configured in an RCST contain primarily information pertaining to that RCST. It may also contain information pertaining to the other RCST involved in the connection, in the case of some particular transparent mesh connections; the corresponding parameters will be prefixed by "Other" (e.g. Other Channel_ID, Other Route_ID, Other Group_ID and Logon_ID).

7.2.3.2 Packet Classification Table

Packet Classification Table (PktClass Table in SatLabs abbreviation) defines the packet classification used in RCST. Packet classification is based on filter criteria/masks including primarily layer 3 (IP) parameters but also some layer 2 Ethernet parameters (e.g. VLAN Priority. A brief description of each parameter is given in table 7.2, which include SatLabs legacy parameters and a set of new parameters required for C2P (shown in fair font) in order to:

- accommodate the IP classes of service specific to IntServ;
- include priority among IP flows;
- include traffic and policy data; and
- include statistics.

A new version of the table, including all legacy and new parameters, will be defined for the purpose of backwards compatibility.

As a result of classification, each packet is associated with a packet type or flow type, referenced by a Packet Classification Index (PktClass Index).

There is one entry in the Packet Classification Table for each defined packet/flow type. A default entry should be defined in order to prevent dropping the packets that do not fall into a classification criteria. A packet that does not match any configured filter, including the default filter, will be discarded.

Each entry in the table is associated with a PHB Entry in the PHB/CoS Mapping Table (clause 7.2.3.3) and a Connection Profile Entry in the Connection Profile Mapping Table (clause 7.2.3.5).

Table 7.2: Packet Classification Table

IP Packet Class Table	No of Bits	Description
PktClassEntry		An entry in the Packet Classification Table
PktClassIndex	16	Index of the Packet Classification Table. Used to identify a packet type or flow type
PktClassDscpLow	32	Specifies the low value of a range of DSCP values to which a packet is compared. A value of 0 is used to inactivate
PktClassDscpHigh	32	Specifies the high value of a range of DSCP values to which a packet is compared. A value of 63 is used to inactivate
PktClassDscpMarkValue	32	Specifies the DSCP value used to mark (remark) a packet. Possible DSCP mark values are (0..63). A value of 64 indicates no DSCP marking
PktClassIPProtocol	32	Specifies the IP protocol to which a packet is compared (e.g. TCP, UDP, etc.)
PktClassIPSrcAddr	32+8/128+8 (see note 1)	Specifies the IP source address to which a packet is compared
PktClassIPSrcAddrMask	32+8/128+8 (see note 1)	Specifies which bits of the IP source address will be matched
PktClassIPDstAddr	32+8/128+8 (see note 1)	Specifies the IP destination address to which a packet is compared. See note 2
PktClassIPDstAddrMask	32/+8/128+8 (see note 1)	Specifies which bits of the IP destination address will be matched
PktClassSrcPortLow	32	Specifies the low range of the source port number to which a packet is compared
PktClassSrcPortHigh	32	Specifies the high range of the source port number to which a packet is compared
PktClassDstPortLow	32	Specifies the low range of the destination port number to which a packet is compared
PktClassDstPortHigh	32	Specifies the high range of the destination port number to which a packet is compared
PktClassVlanPri	32	Specifies the VLAN User Priority to which a packet is compared
PktClassPhbAssociation	8	Associates the filter entry to a specific PHB/IntServ CoS (by reference to a PhbIndex in the PHB/CoS Mapping Table)
PktClassIntServCoSId	1	Specifies the identifier of the IntServ Class of Service of the packet, as obtained from a control interface. Possible values: "0":GS; "1": CL (see note 4)
PktClassIPPriority	8	Applicable to IP elementary flows (e.g. IntServ), for priority scheduling (0: No priority; 1: Low priority; 255: High priority). See note 5
PktClassTraficSpecPolicyData	256 (see note 3)	Relevant traffic information/policies (typically derived from a control plane interface
PktClassCnxProfileAssociation	16	Associates this entry to a specific connection profile (by reference to a CnxProfileIndex in the Connection Profile Mapping Table)
PktClassAction	1	Specifies if the packets mapped to this entry (flow type) can be transmitted to the satellite interface or should be discarded ("0": Permit; "1": Deny). The parameter can be related to a firewall function, used to avoid undesired incoming traffic
PktClassOutOctets	32	Number of octets sent out over the satellite
PktClassOutPkts	32	Number of packets sent out over the satellite (can be unicast or multicast, depending on the packet classification)
PktClassRowStatus	8	Standard SNMP row status

IP Packet Class Table	No of Bits	Description
NOTE 1: The present specification supports both IPv4 and IPv6 addressing. IPv4 allocates 32 bit (+8 bit mask) addresses and IPv6 allocates 128 bit addresses (+8 bit mask).		
NOTE 2: The PktClassIPDstAddr can contain an IP subnet that corresponds to several peer parties (e.g. several peer RCSTs).		
NOTE 3: The size of this field is a rough estimate. It will depend on the specific form of traffic information/policies (implementation specific).		
NOTE 4: Two classes of service are defined in IntServ - Guaranteed Service (GS) and Control Load (CL). The class of service of an IP flow is obtained from an RSVP message.		
NOTE 5: Used for internal processing in RCST (e.g. IP layer priority scheduling).		

Traffic information/policies in table 7.3 typically include specifications related to the traffic characteristics/level, service requirements/desired QoS, sender/receiver QoS handling capabilities, etc. Figure 7.4 provides an example of IP policy, as defined in the IntServ RSVP message parameters (RFC 2205 [i.19]). Similar policies/parameters can be defined for DiffServ, in association with the use of and end-to-end signalling protocol (e.g. SIP [i.20]/SDP [i.22]) or based on the Service Level Agreement (SLA)/Traffic Conditioning Agreement (TCA).

Table 7.3: Traffic Policies/Parameters for IntServ RSVP Messages

Message	RSVP Object	IntServ parameters	Description
RESV	FILTER SPEC		Sender IP address, optionally UDP/TCP port number
	FLOW SPEC	TSPEC	Level of traffic
		RSPEC	Bandwidth delay
		Desired QoS	GS, CL services
PATH	SENDER_TEMPLATE		Sender IP address, optionally the UDP/TCP sender port, protocol ID for the session
	SENDER_TSPEC		Describes the traffic the application expects to generate
	ADSPEC		QoS control capabilities and requirements of the sending application (updated by network nodes)

The RCST should be capable of handling a number of flow types predefined for anticipated applications and priorities. The flow types defined in a system are system-specific. Table 7.4 provides an example of C2P related parameters for a number of flow types defined for some popular applications. The parameters are system specific.

Table 7.4: Flow Types and C2P Related Parameters

Flow Type	Application	IP CoS/PHB	Uni/Bi	SDR return (kbps)	PDR return (kbps)	SDR forward (kbps)	PDR forward (kbps)
Type 1	VoIP (aggregate)	EF	Bi	16	16	16	16
Type 2	Videoconference	EF	Bi	384	384	384	384
Type 3	Interactive voice/video	AF1	Bi	256	384	256	384
Type 4	Critical data	AF2	Bi	32	48	16	16
Type 5	Call signalling	AF3	Uni/Bi	16	32	-/16	-/16
Type 5	Data streaming	AF4	Uni	64	128	-	-
Type 6	FTP, Email	BE	Bi	16	16	16	256
Type 7	Web browsing	BE	Bi	16	16	64	256
Type 8	Audio-video playback	GS	Uni	16	16	-	-
Type 9	Instant messaging	CL	Uni	48	64	-	-
Type 10	Default	BE	Bi	16	16	16	16

NOTE 1: Associating a given PHB with different applications (e.g. EF and BE in the table) indicates that the same service (packet forwarding) is expected for all applications. The differentiation between flows can be achieved by using different DSCPs, associated with Class Selector (CS) PHBs with forwarding rules similar to those of the indicated PHB.
NOTE 2: For some applications/flow types the table may include additional parameters, such as Priority, Traffic specification, Service specification (delay, jitter), etc.
NOTE 3: The connection associated with Type 5 can be unidirectional or bidirectional, depending on the type of signalling.
NOTE 4: For Type 10 the rate parameters can be zero if a connection for BE has already been established.

In the case of flow types pointing to a bidirectional connection, it is assumed that the same QoS architecture, the same classification criteria/masks and the same flow types are configured at both ends of the connection.

The rate parameters identified in the above table, with values as defined by the system/network operator, can be used for the creation of connection profiles in the Connection Profile Mapping Table (clause 7.2.3.3) for the flow types/filter patterns defined for the system.

7.2.3.3 Connection Profile Mapping Table

Connection Profile Mapping Table includes the set of parameters defining the connection (or C2P) profile (as per clause 5.4.4). They are used to specify the profile-related parameters in the C2P connection establishment/modify messages.

Connection Profile Mapping Table will be configured in each RCST by management or a local interface (e.g. CLI or web interface).

The minimum information required per connection profile entry, uniquely identified by a Connection Profile Index (CnxProfileIndex), is included in table 7.6.

Table 7.5: Connection Profile Mapping Table

Connection Profile Table	No of Bits	Description
CnxProfileEntry		An entry in the Connection Profile Mapping Table
CnxProfileIndex	32	Index (identifier) of the Connection Profile
CnxProfileConnectionType	8	Defines the type of the connection for the IP flow/aggregate (unicast or multicast, unidirectional or bidirectional, RCST/RSGW-initiated or NCC-initiated) mapped to this connection profile
CnxProfileC2PCoS	8	MAC class of service for the connection profile (assumed the same in both directions in the case of a bidirectional connection), see note.
CnxProfileReturnSDR	8	Transmit sustainable data rate for the connection, required for the transmission of packets corresponding to an IP flow of the type mapped to this entry
CnxProfileReturnPDR	8	Transmit peak data rate for the connection, required for the transmission of packets corresponding to an IP flow of the type mapped to this entry
CnxProfileForwardSDR	8	Receive sustainable data rate for the bidirectional connection, required for the reception of packets corresponding to an IP flow of the type mapped to this entry
CnxProfileForwardPDR	8	Receive peak data rate for the bidirectional connection, required for the reception of packets corresponding to an IP flow of the type mapped to this entry
CnxProfileInactivityTimeout	16	This parameter is used to trigger the release of the connection if there is no forward or return traffic activity (in seconds; 0 = never times out)
CnxProfileRowStatus	8	Standard SNMP row status
NOTE: In the case of bidirectional connection the same C2P CoS assumed for both directions implies that identical or similar QoS architectures are implemented at both ends of the connection.		

Each entry in the table defines the connection profile for a given flow type. The number of entries in the table is equal to the number of flow types defined for the system. A number of connection profiles could be predefined for the anticipated applications/flow types (see clause 7.2.3.2).

The SDR and PDR parameters in the Connection Profile Mapping Table (Return SDR/PDR and also the Forward SDR and PDR in the case of bidirectional connections) offer a convenient mechanism to specify the bandwidth resources that need to be requested for a connection in order to carry a flow of the type mapped to a specific connection profile entry, when the connection messages for connection establishment/modify are triggered by events in the user plane, such as the arrival of an IP packet. In the case where the connection establishment/modify messages were triggered by events in the control plane, e.g. the interception of an application session signalling message, the bandwidth parameters could be provided by the corresponding session signalling protocols (e.g. RSVP, SIP/SDP).

7.2.3.4 Active Connection Table

The Active Connections Table contains all data necessary to fully characterize all active (open) connections of a given RCST. The table includes an entry for each active connection. The entry is dynamically created and updated every time a C2P message is received by the terminal.

The minimum information required per active connection, uniquely identified by a Connection Reference ID (ActiveCnxRefId), is included in table 7.7. Each entry in the table is associated to a Request Class Entry in the Request Class Table and to a Packet Class Entry in the Packet Classification Table.

The majority of the parameters in the table (e.g. addresses, VPI/VCI, PID, Channel_ID, Route_ID) are configured when the active connection is created. The rate parameters of the connection (Return SDR/PDR and also Forward SDR/PDR in the case of bidirectional connections) can be dynamically updated by connection modify request/response messages. The Active Connection Table includes both requested rate parameters and admitted rate parameters.

The requested rate parameters (ReqReturnSDR/PDR, ReqForwardSDR/PDR) represent the requested aggregate rates for all IP flows/flow aggregates mapped to a specific active connection at a given time. They are included in all C2P request messages related to the connection. The first request is for connection establishment, triggered e.g. by the first flow that is mapped to this connection, and the requested rates are obtained from the Connection Profile Mapping Table. The rates are then updated with each flow/flow aggregate that is activated/deactivated within the connection, and a C2P request for a connection modification is sent. The request includes the total (cumulative) updated rates, not just the incremental changes.

The admitted rate parameters (AdmReturnSDR/PDR, AdmForwardSDR/PDR) are dynamically updated by the NCC via C2P messages, in response to each C2P connection establishment/modify request, after the execution of a connection admission control function (CAC).

In the case of unidirectional connections, the rate parameters apply to one direction and the table includes only one set of requested/admitted rate parameters - either the ReturnSDR/PDR (if the connection was initiated by the RCST) or the ForwardSDR/PDR (if the connection was initiated by the other RCST involved in the connection or the NCC).

In the case of bidirectional connections, the requested/admitted rate parameters will be updated for both directions, regardless of the initiator of the connection establishment/modify request.

In the case of bidirectional connections between two RCSTs, identical or similar parameters (with the same terminology) will be maintained in the Active Connection Tables at the two ends of the connection (for a given ActiveCnxEntry), but the meaning is changed: the Source, Return and Forward parameters at one RCST become Destination, Forward and Return parameters, respectively, at the other RCST. For example, the requested/admitted Return and Forward SDR/PDR at one RCST will be reversed at the other RCST, and vice-versa.

With regard to the C2P CoS, Connection Type and Connection Inactivity Timeout, they will be identical at the two ends of the connections, which are assumed configured to support the same QoS architecture/same PHBs and/or IP CoSs. The Connection Inactivity Timeout is set by the NCC and may be different from that in the Connection Profile Mapping Table, used in the connection establishment request message.

Table 7.6: Active Connection Table

Active Connection Table	No of Bits	Description
ActiveCnxEntry		An entry in the Active Connection Table
ActiveCnxIndex	32	Index of the active connection
ActiveCnxRefId	16	Identifier of the active connection
ActiveCnxMACSrcAddr	48	MAC address of the RCST (see note 2)
ActiveCnxCause	8	Last reported cause of the active connection
ActiveCnxChannelId	8	Channel_ID of the connection; same value as in RC table
ActiveCnxMACDestAddr	48	MAC address of the peer RCST
ActiveCnxReturnStreamId	24	Transmit VCC for ATM or transmit PID for MPEG
ActiveCnxForwardStreamId	24	Receive VCC for ATM or receive PID for MPEG
ActiveCnxType	8	Defines the type of the connection (unicast or multicast, unidirectional or bidirectional, RCST/RSGW-initiated or NCC-initiated)
ActiveCnxC2PCoS	8	MAC class of service for the connection (assumed the same in both directions in the case of a bidirectional connection). Example of values (SatLabs): "1": Real Time (RT), "2": Critical Data (CD), "3": Best Effort (BE) and "4":Network Management (NM)
ActiveCnxReqReturnSDR	8	Requested transmit sustainable data rate of all flows mapped to the connection
ActiveCnxReqReturnPDR	8	Requested transmit peak data rate of all flows mapped to the connection
ActiveCnxReqForwardSDR	8	Requested receive sustainable data rate of all flows mapped to the connection. (Informational)
ActiveCnxReqForwardPDR	8	Requested receive peak data rate of all flows mapped to the connection. (Informational)
ActiveCnxAdmReturnSDR	8	Admitted transmit sustainable data rate of all transmitted flows carried by the connection
ActiveCnxAdmReturnPDR	8	Admitted transmit peak data rate of all transmitted flows carried by the connection
ActiveCnxAdmForwardSDR	8	Admitted receive sustainable data rate of all received flows carried by the connection. (Informational)
ActiveCnxAdmForwardPDR	8	Admitted receive peak data rate of all received flows carried by the connection. (Informational)
ActiveCnxRouteId	16	Route_ID associated to the connection; equivalent to a Channel_ID list
ActiveCnxGroupLogonId	24	Group_ID and Logon_ID of the RCST
ActiveCnxOtherChannelId	8	Channel_ID of the other RCST (see notes 1 and 2)
ActiveCnxOtherRouteId	16	Route_ID of the other RCST (see notes 1 and 2)
ActiveCnxOtherGroupLogonId	24	Group_ID and Logon_ID of the other RCST (see notes 1 and 2)
ActiveCnxInactivityTimeout	16	Used to trigger the release of the connection if there is no forward or return traffic activity (in seconds; 0 = never times out)
ActiveCnxMaxPacketSize	16	Maximum packet size
ActiveCnxStatus	8	Status of the connection: 1:SetupInProgress, 2:ConnectionOpen, 3:ConnectionReleaseInProgress, 4:ConnectionModifyInProgress
ActiveCnxPktClassAssociation	32	Associates this entry to a specific Packet Class (by reference to a PktClassIndex in the Packet Classification Table)
ActiveCnxRequestClassAssociation	32	Associates this entry to a specific request class (by reference to a requestClassIndex in the Request Class Table)
ActiveCnxOutOctets	32	Number of octets sent out over the satellite
ActiveCnxOutPkts	32	Number of packets sent out over the satellite (can be unicast or multicast, depending on the packet classification)
ActiveCnxInOctets	32	Number of octets received from the satellite
ActiveCnxInPkts	32	Number of packets received from the satellite (can be unicast or multicast)
List of ActiveCnxIPv4SrcSubnets	32+8	Defines the IPv4 source addresses and masks for all subnets
List of ActiveCnxIPv4DstSubnets	32+8	Defines the IPv4 destination addresses and masks for all subnets
List of ActiveCnxIPv6SrcSubnets	32+8	Defines the IPv6 source addresses and masks for all subnets
List of ActiveCnxIPv6DstSubnets	32+8	Defines the IPv6 destination addresses and masks for all subnets
ActiveCnxRowStatus	8	Standard SNMP row status
NOTE 1: This parameter may be needed for bidirectional connections defined for some particular transparent mesh scenarios. The prefix "Other" is used for the RCST at the other end of the connection, referenced with regard to the RCST to which this table applies.		
NOTE 2: Not used in the C2P messages defined in this version of the specification.		

7.2.3.5 PHB/CoS Mapping Table

The PHB/CoS Mapping Table is a SatLabs legacy table, extended to satisfy additional C2P needs.

The original SatLabs table, labelled PHB Mapping Table, was changed to PHB/CoS Mapping Table and exrtended in order to reflect the inclusion of the following parameters:

- IP classes of service (IP CoS) specific to IntServ;
- parameters specific to each PHB/CoS (e.g. rates, delay, jitter, buffer size);
- parameters associated with a defined mechanism, e.g. a token bucket mechanism (optional); and
- statistics.

The parameters to be included in the PHB/CoS Table are shown in table 7.8. The new parameters and the changes to the original parameters are shown in fair font. There is one entry in the table for each PHB/CoS. Each entry is associated with a Request Class Entry (requestClassEntry) in the Request Class Table. For simplicity all parameters include the PHB prefix, even if they may refer to an IntServ CoS (i.e. GS or CL).

Table 7.7: PHB/CoS Mapping Table

PHB/CoS Table	No. of bits	Description
PhbMappingEntry		An entry in the PHB mapping table
PhbIdentifier	32	Identification of the PHB in the range (0 to 63)
PhbName	String	Name of the PHB
PhbIPCoSIdentifier	8	Identification of the IntServ CoS in the range (0, 2) (0: not valid; 1: GS; 2: CL), see note 3
PhbIPCoSName	String	Name of the IntServ CoS
PhbAdmReturnSDR	8	Admitted (true) transmit sustainable data rate (in bits/s), see note 1
PhbAdmReturnPDR	8	Admitted (true) transmit peak data rate (in bits/s), see note 1
PhbBufferSize	16	Traffic buffer size (in bytes)
PhbDelay	16	Delay specification (in milliseconds)
PhbJitter	16	Jitter specification (in milliseconds)
PhbPolicy	128 (see note 2)	As per DiffServ standard or system specific
PhbTokenBucketR1	16	First rate of a two-rate token bucket mechanism (in bits/s), (optional)
PhbTokenBbucketR2	16	Second rate of a two-rate token bucket mechanism (in bits/s), (optional)
PhbTokenBucketSize	16	Token buffer size (in bytes), (optional)
PhbRequestOutOctets	32	Number of octets sent out over the satellite
PhbRequestOutPkts	32	Number of packets sent out over the satellite (can be unicast or multicast, it depends of the packet classification)
PhbRequestOutDiscards	32	Number of outbound packets that were free of errors but discarded. (i.e. packets that were filtered out, e.g. to free up memory)
PhbRequestOutErrors	32	Number of outbound packets discarded because of errors.
PhbRequestOutQLen	32	Number of packets in the outbound queue
PhbRequestClassAssociation		This object provides an association of this PHB/IP CoS to a Request Class (by referencing to a RequestClassIndex in the Request Class Table)
PhbMappingRowStatus	8	Standard SNMP row status
NOTE 1: This parameter may be used for internal processing in RCST (traffic conditioning), e.g. to derive the optional token bucket rate parameters. NOTE 2: The size of this field is a rough estimate; it will depend on the specific form in which the PHB policies are specified (implementation specific). The PHB policy in this table should be consistent with the IP traffic specification and policy data in the Packet Classification Table. NOTE 3: Two classes of service are defined in IntServ - Guaranteed Service (GS) and Control Load (CL). The class of service of an IP flow is obtained from an RSVP message.		

The PHB/CoS table is configured by management. The rate parameters (including the token bucket rates) are dynamically updated as connections are established/modified, by relying on the Active Connection association with a Packet Class and the Packet Class association with a PHB/CoS. For a given PHB/CoS the admitted rate parameters AdmReturnSDR/AdmReturnPDR should reflect the components (that can be associated with this PHB/CoS) of the AdmReturnSDR/AdmReturnPDR parameters of all connections to which this PHB/CoS is mapped. The calculation of the PHB/CoS rate parameters from the connections' rate parameters is performed by each RCST based on rules that are system-specific. The derivation of the token bucket rate parameters from the AdmReturnSDR and AdmReturnPDR is based on rules that are both PHB/CoS-specific and implementation specific.

NOTE: According to the cardinality model (clause 6.1.3):

- A PHB/CoS can be mapped to multiple connections.
- Several PHBs/CoSs can be mapped to one connection (i.e. flows with different PHB/CoS can share one connection), therefore only a fraction of connection's rate can be attributed to a given PHB/CoS.

7.2.3.6 Request Class Table

The Request Class Table is a SatLabs legacy table. It is slightly changed to include some C2P-specific parameters, namely:

- Route_ID;
- statistics.

The Request Class Table defines all layer 2 QoS parameters for each supported RC, i.e. the MAC layer logical parameters (Channel_ID, PIDs or VPI/VCI) and bandwidth parameters, expressed as capacity categories limit values (CRA, RBDCMax, VBDCMax). These parameters are controlled by the NCC and configured in both NCC and RCST. In RCST they are used for capacity request calculation and possibly in the packet scheduling/dispatching process. In NCC they are used for capacity assignment in response to dynamic requests from RCSTs.

The Request Class Table is defined in 7.2. It has one entry for each RC, identified by a Request Class Index (RC Index). The new parameters are shown in fair font.

Table 7.8: Request Class Table

Request Class Table	No. of bits	Description
requestClassEntry		An entry in the Request Class (RC) table
requestClassIndex	32	Index of the Request Class table. A maximum of 256 entries can be created
requestClassName	String	Name of the RC
requestClassChanID	32	Channel_ID of this RC
requestClassVccVpi	32	VPI used for the RC (ATM profile)
requestClassVccVci	32	VCI used for the RC (ATM profile)
requestClassPidPoolReference	32	Reference to the preferred PID in the PID pool applicable to the RC (typically 1 PID per RC is used). PID pool table is as defined in [i.4]
requestClassCRA	32	Defines the CRA level for the Request Class
requestClassRBDCMax	32	Max RBDC that can be request for the Request Class
requestClassRBDCTimeout	32	Persistence of the RBDC request
requestClassVBDCMax	32	Max VBDC that can be allocated for the Request Class
requestClassVBDCTimeout	32	Time after which the RCST considers that the pending request are lost
requestClassVBDCMaxBackLog	32	VBDC Backlog per Request Class
requestClassRouteId	16	Route_ID associated to the request class; equivalent to a Channel_ID list
requestClassOutOctets	32	Number of octets sent out over the satellite (including the ATM or MPEG header)
requestClassOutCells	32	Number of cells (ATM or MPEG) sent over the satellite
requestClassRowStatus	8	Standard SNMP row status

The Request Class Table is configured by management with initial values and then the capacity limit values (CRA, RBDCMax, VBDCMax) can be dynamically updated (based on C2P messages) in response to a specific C2P RC modify request or when a connection is created, modified or released. In the latter case (which corresponds to an implicit RC modify request), the per-RC capacity limit values should reflect the connection admitted transmit parameters (i.e. AdmReturnSDR/PDR, as captured in the Active Connection Table) of all connections mapped to this RC. The conversion of the aggregate (per-RC) rate parameters to capacity limit values is performed by the NCC, based on rules that are both RC-specific and system-specific. The per-RC capacity limit values should also be consistent with the return rate parameters of all CoSs/PHBs mapped to an RC.

The VCC/VPI and VCC/VCI in the Request Class Table correspond to the Transmit VCC (as Return Stream ID) in the Active Connection Table. In case of conflict, the VCC in the Active Connection Table takes precedence over the VCC in the Request Class Table.

The PID (as Return Stream ID) in the Active Connection Table corresponds to the preferred PID for the Request Class in the PID Pool referenced in the Request Class Table. If only one PID per Request Class is defined, it is the same as the PID specified in the Active Connection Table. In case of conflict the PID in the Active Connection Table takes precedence over the PID in the Request Class Table.

7.2.3.7 Table Linkage Mechanism for Dynamic Connectivity/RCST Behaviour

The derivation of C2P-related and QoS-related parameters associated with the dynamic connectivity is based on linking various tables defined as part of the RCST data structure. The table linkage mechanism is illustrated in figure 7.3. The links between various tables and the C2P request/response are numbered in order to facilitate the description.

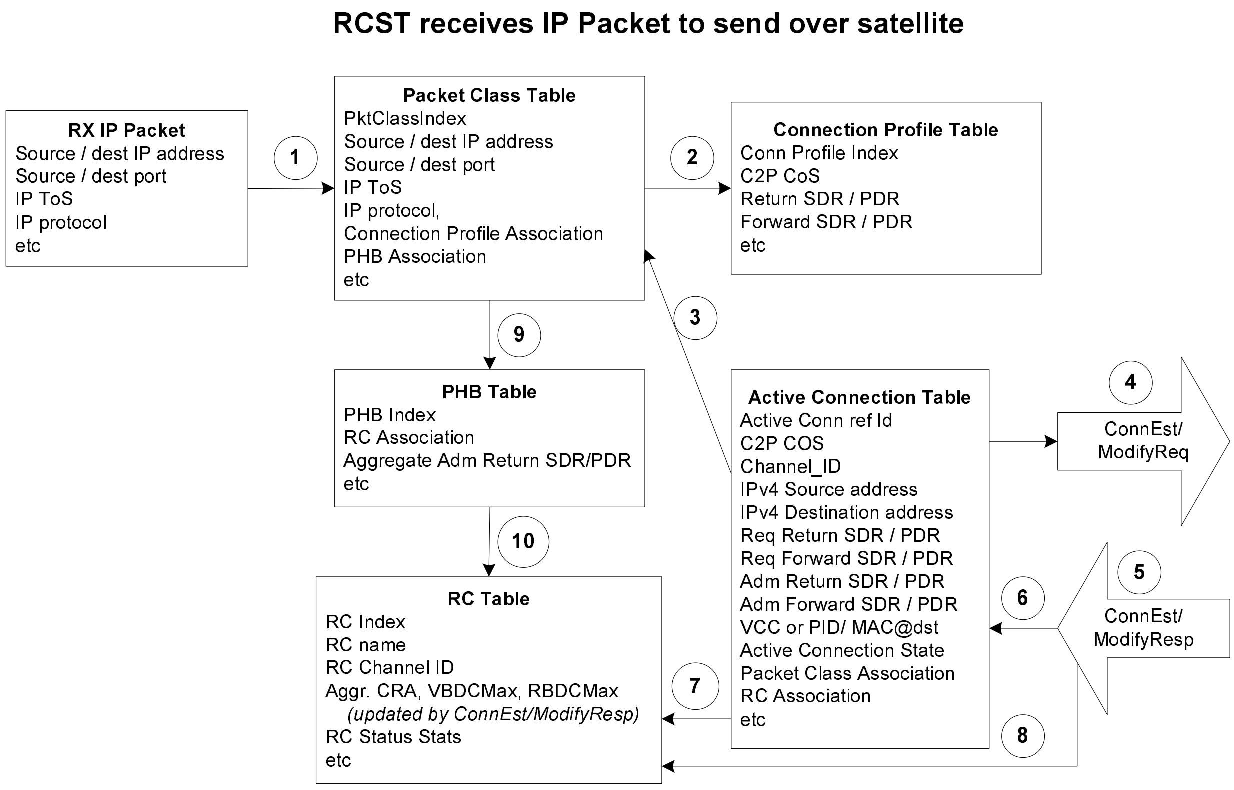

Figure 7.3. RCST Data Structures Linkage/RCST Behaviour

The RCST detects (1) an incoming IP packet, to be forwarded to the satellite interface. The packet is classified based on a combination of parameters included in the Packet Classification Table, used to define filter masks. If the information in the packet header matches one of the existing (configured) filter masks, the packet is classified into a pre-defined packet type or flow type (including a default flow type), identified by a PktClassIndex; if not, the packet is discarded. The flow type applies to either an elementary IP flow (IntServ) or to an aggregation of IP flows/PHB (DiffServ).

The flow type with its assigned PktClassIndex is associated (2) with a connection profile in the Connection Profile Mapping Table, including the nature of the connection (unidirectional or bidirectional) and a number of C2P relevant parameters - in particular the IP CoS/PHB, the return SDR and PDR and also the forward SDR and PDR (if the connection is bidirectional).

The flow type/PktClassIndex of the incoming packet is also associated (9) with a PHB/IP CoS in the PHB/CoS Mapping Table, which is in turn associated (10) with a Request Class in The Request Class Table.

Once a packet has been classified to PktClassIndex, the RCST checks whether an active connection (i.e. an entry in the Active Connection Table) exists for the PktClassIndex and the IP destination address of the incoming IP packet. Depending on the outcome, the following actions can be taken:

- If no entry is found, an entry (ActiveCnxEntry) will automatically be created in the Active Connection Table, with the profile taken from the Connection Profile Mapping Table, and a Connection Establishment Request will be sent (4).
- If one or more entries matching the PktClassIndex are found but none matches the IP destination address of the incoming IP packet, an entry will be created in the Active Connection Table, with the profile taken from the Connection Profile Mapping Table, and a Connection Establishment Request will be sent (4).
- If an entry is found for the PktClassIndex but the packet is not part of an existing IP flow/flow aggregate, the request bandwidth of the corresponding connection will be updated in the Active Connection Table and used in a subsequent connection modify request (4), issued in order to accommodate the new flow.
- If an entry is found and the packet is already part of an existing IP flow/flow aggregate carried by the connection (as identified by the PktClassIndex), the packet will be forwarded according to the parameters stored in the Active Connection Table.

NOTE 1: A flow type, identified by a PktClassIndex, is defined for a destination IP subnet and may correspond to several RCSTs.

NOTE 2: An active connection is created for a specific destination IP address. Several connections can be created for a given PktClassIndex, for different destination IP addresses, all part of the same destination IP subnet.

NOTE 3: An active connection carries one or multiple flow(s) or flow aggregate(s) of the same type.

The NCC response message (5) will carry the connection parameters or updates of the connection parameters (admitted return and forward SDR and PDR), which will be used to update the Active Connection Table (6).

The NCC response also carries updates of the capacity limit values (CRA, RBDCMax, VBDCMax) of the Request Class to which the connection is mapped, as given by the RC association in the Active Connection Table (7). They will be used to update the Request Class Table (8). The RC capacity parameters are a reflection of the rate parameters of all connections mapped to the RC (clause 7.2.3.6).

The rate parameters in the PHB/CoS Mapping Table (i.e. the admitted return SDR/PDR and also the token bucket rates) are dynamically updated with each connection establishment/modify response, by relying on the Active Connection association with a Packet Class and the Packet Class association with a PHB/CoS. For a given PHB/CoS, the admitted rate parameters AdmReturnSDR/AdmReturnPDR are a reflection of the components (that can be associated with this PHB/CoS) of the AdmReturnSDR/AdmReturnPDR parameters of all connections to which this PHB/CoS is mapped (clause 7.2.3.5).

NOTE 4: Each IP flow/flow aggregate is mapped to only one PHB/CoS and to only one active connection.

NOTE 5: Each active connection is mapped to a Request Class.

NOTE 6: Several PHBs/CoSs can be mapped to a Request Class.

7.3 IP Addressing and Routing

DVB-RCS networks can be seen as IP networks, meaning that IP packets are routed through the system according to the destination IP address contained in the header and the routing information configured in the network nodes, in particular in the ground equipment (RCST, GW, RSGW). From a network viewpoint the RCST itself behaves as an IP router.

In clause 7.2, it was shown that C2P supports address resolution for the purpose of MPE encapsulation (i.e. the association between MAC addresses and IPv4 addresses of both parties involved in the connection). This information is provided at connection setup in the C2P response message from the NCC to the RCSTs.

C2P allows to dynamically update the routing information (entries in the routing table). The RCST will maintain the routing data in the Active Connection Table until the connection is released.

C2P's embedded address resolution function allows the NCC to be seen as a centralized router. It knows all the IP subnetworks in the satellite network. The RCST no longer needs to be configured with all possible static IP routes. The C2P messages exchanged by the RCST may dynamically update its routing information, as part of the Active Connection Table maintenance.

The routing information is based on the Classless Inter-Domain Routing (CIDR) notation (RFC 1518 [i.36] and RFC 1519 [i.37]). The routing table in the RCST is typically looked up by applying the best match rule to select the connection with the longest matching prefix (i.e. smallest subnet). Each entry in the RCST routing table is associated with a destination IP address which can identify:

- a network prefix route: i.e. the stored IP address has a prefix length between 1 and 31 bits (127 in case of IPv6);
- a host-specific route: i.e. the stored IP address has a prefix length of 32 bits (128 in case of IPv6);
- a default route (prefix length of 0 bits).

NOTE: The address coding in CIDR notation uses the 5 byte format aa.bb.cc.dd/ee (4 bytes for IPv4 address + 1 byte for the shortened mask value). The mask value (or length) defines the number of left-most contiguous mask bits that are set to one.

The process is illustrated in figure 7.4. When an IP packet is received at the RCST's User Interface (UI), the destination address in the IP header is compared with the target IP addresses stored in the routing table (static routing entries) (1). The selected routing table entry is that associated with the destination address presenting the longest prefix length, matching the destination address of the incoming IP packet. If no match is found the packet is dropped (2). If a match is found (3), the RCST has to decide whether this packet is to be forwarded to the satellite interface (5) or towards the LAN interface (4). This decision is taken thanks to the routing configuration in the RCST routing table (i.e. static routing entries). If the packet is to be forwarded towards the satellite interface, the RCST should first classify the packet according the process already described in clause 7.2.3.7.

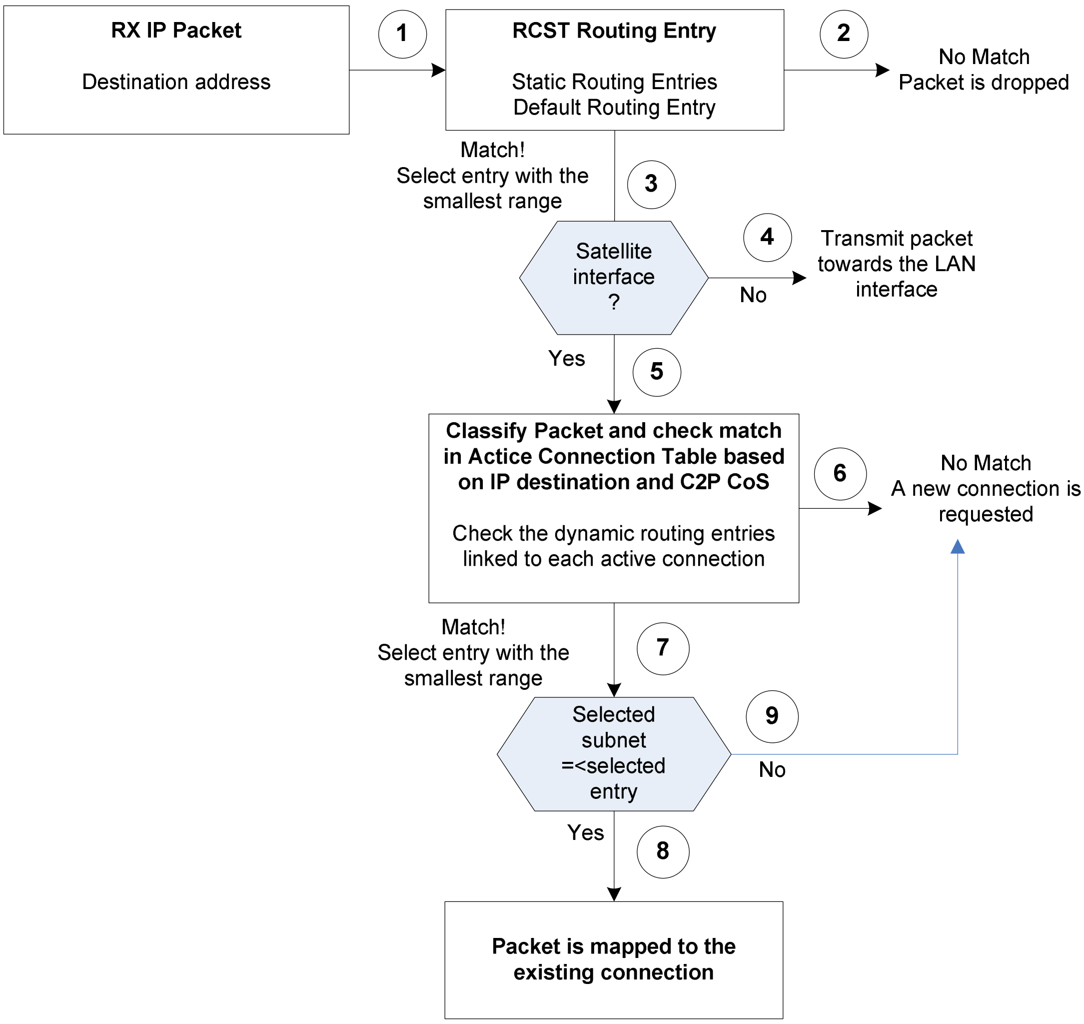

Figure 7.4. IP Routing and New Connection Triggering

Once a packet has been classified to a PktClassIndex, the RCST checks whether an active connection (i.e. an entry in the Active Connection Table) exists for the PktClassIndex and the IP destination address of the incoming IP packet. At this point, the RCST will checked the dynamic routing entries associated to each one of the active connections, selecting the entry with the smallest address range.

7.4 Multicasting

The purpose of this clause is to identify various options for dynamic multicast group management, multicast address resolution and multicast tree establishment, and to determine the interaction between these multicast functions and C2P functionalities.

7.4.1 Overview

IP Multicast delivers source traffic to multiple receivers without adding any additional burden on the source or the receivers, while using the least network bandwidth of any competing technology.

As already identified in i.9, multicast operations over BSM networks (and thus over DVB-RCS networks) can have multiple forms involving different attributes:

- IP multicast group management attributes:
  - static: the terminal statically forwards all received groups or a pre-configured set of groups; and
  - dynamic: the terminal allows the attached hosts to dynamically set which groups to forward.

- Multicast address resolution:
  - static: all multicast addresses are preconfigured and only refreshed infrequently via tables or direct operator's intervention; and
  - dynamic: multicast addresses are mapped dynamically using an address resolution protocol.
- Multicast tree establishment over satellite based on the multicast group management.

Scheduled and static operations do not necessitate any IP layer protocol operations beyond the usual connectivity performed based on C2P.

For dynamic multicast management at the IP layer where hosts join and leave multicast sessions "at any time", the following essential aspects need to be covered:

1) IP Multicast group management via the Internet Group Management Protocol (IGMP) [i.25] proxying or snooping and protocol adaptation that support dynamic join and leave operations. This is based on IGMPv2 protocol satellite adaptation already defined in [i.9].

2) The use of IP Multicast routing protocols such as the PIM-SM [i.24] but not commonly used in DVB-RCS networks.

3) Multicast address resolution, to resolve an IP multicast address to a stream ID. This is solved in DVB-RCS networks by using the Multicast Map Table (MMT). This table provides the mapping of multicast groups to MPEG2 Packet Identifiers (PIDs). The MMT format is currently included in SatLabs specifications and it is expected to be included in the next version (V1.5.1) of DVB-RCS standard.

4) Multicast Security, to define how security policies can be applied to multicast transmissions (e.g. only authorized users can receive the content and that the content is protected by the some form of Digital Rights Management [DRM] or merely by the definition of multicast address scoping).

IP multicast protocols and their signalling characteristics were originally developed on terrestrial wired LAN or WAN technologies.

An additional consideration is that multicast implementation in any DVB-RCS network interconnection may involve replication to reach all intended destinations. The integration of a satellite network with other networks raises some issues concerning replication. Replication outside the satellite network may not optimize satellite resource usage. As for replication within the satellite network, the typical scenario in transparent satellite systems is replication at the gateway, with one copy for each beam with a receiving terminal. With OBP satellites on board replication is possible, yielding the most efficient spectrum usage. Finally, destination terminals may replicate data depending on how their local network handles multicast.

The above overview helps to introduce the possible static/dynamic multicast scenarios based on a hybrid solution that combines the IGMP, adapted to satellite environment, and the C2P. The general architecture used for IGMP adaptation follows the multicast reference architecture defined in [i.10].

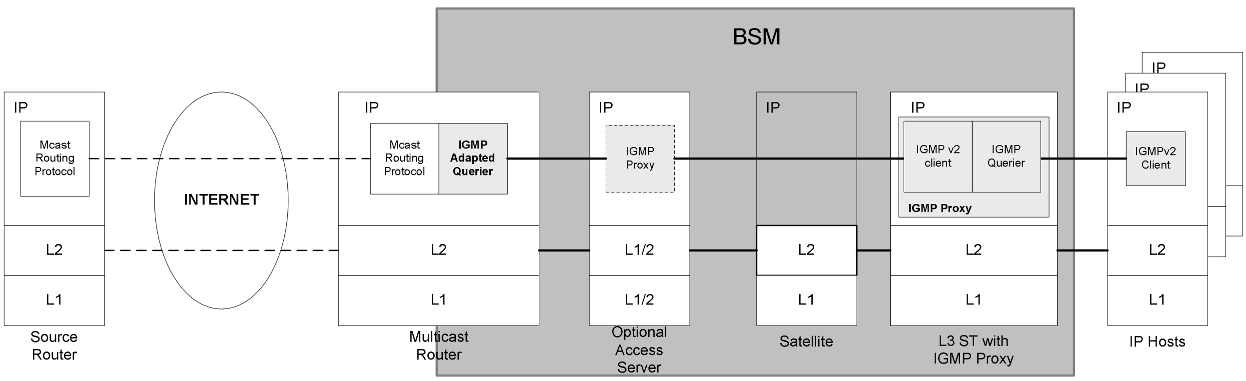

Figure 7.5: General Multicast Architecture Based on IGMP with Layer 3 ST (routing) with Optional IGMP Proxy

7.4.2 IGMPv2 Satellite Adaptation

NOTE: This clause is only applicable to IGMPv2 [i.25] or IGMPv3 [i.26] running in IGMPv2 compatibility mode.

The Internet Group Management Protocol is an integral part of IP and is required to be implemented by all hosts wishing to receive IP multicast.

IGMPv2 is well suited for a terrestrial network with a shared medium. Indeed every host can listen to the reports transmitted by other hosts. If a report is already transmitted, the host will also stop the ongoing timer and will not send its report. This technique will prevent hosts from replying to useless reports.

The situation is different in a satellite context. Satellite terminals can not directly listen to the replies from other member terminals. Moreover, as the satellite multicast groups can be very large (e.g. satellite networks with 5 000 RCSTs) and very dynamic, this issue can have some serious consequences: the IGMP querier located behind the NCC/GW (or RSGW) may receive thousand of reports from RCSTs. This is the "IGMP flooding" issue that causes a waste of bandwidth and a high CPU activity at the querier side.

Another essential issue mostly due to the IGMP behaviour is the latency when the last end user leaves a multicast group. When the last member is leaving the IGMP, implemented as defined in [i.25], will take a long time to detect that there is no active member on the group any longer, before it stops transmitting the multicast stream on the air interface to that group. The time required is defined by the IGMP protocol and will lead to an important waste of bandwidth. IGMP maximum response time is set according to the estimated number of members in the group, but it should be tuned to ensure that few reports are received if there are still active members and to reduce the (leave) latency if there is only one member remaining.

The ETSI specification for IGMP adaptation for Multicast Group Management [i.9] covers the necessary modifications to IGMP, which improve its performance when applied over satellite links. These modifications are internal to IGMP and do not affect the system's interoperability with other IP networks. The modifications concern particularly the so-called "IGMPv2 mode" described in [i.25] which falls under the umbrella of IGMPv3 [i.26].The adapted mode of IGMP is termed S-IGMP.

This solution is suitable for multicast groups on geostationary satellite systems with two-way links. It is suitable for systems based on Any-Source Multicast (e.g. star networks) and especially for those which do not automatically retransmit IP multicast packets on return links to all group members. It is also beneficial to systems which do have this ability for full subnet multicast.

7.4.3 C2P - IP Multicasting by IGMP

The main issues for IP multicasting and C2P interworking are:

1) Multicast connection establishment and satellite channel establishment.

2) Multicast tree establishment over satellite based on the multicast group management.

The first issue could be easily resolved by C2P that allows signalling between the NCC and terminal. The NCC will assign the required resources based on what kind of channel is needed, i.e. one-to-one channel or one-to-many channel. A channel is understood in this context as a combination of connectivity channel(s) and connection(s) as defined for the purpose of C2P (clause 5).

The multicast tree establishment over satellite will depend on the multicast routing protocol used to set up the resources for the links between two or more terminals. Therefore, one has to consider the interworking between its satellite channel establishment procedure and the multicast routing protocol. An OBP satellite provides packet duplication on board and can therefore be treated, together with the NCC, as a multicast enabled forwarder. S-IGMP could be one solution, but also the C2P.

C2P is designed for satellite channel establishment but it can be extended to support a dynamic multicast environment. The suggestion is to limit the C2P to a satellite channel establishment function and add possible extensions for interworking with the multicast routing protocols in such a way that the NCC can always be informed of the multicast groups and of the initiated multicast connections, so it can configure accordingly the OBP if necessary.

The architecture for C2P interworking with IP multicasting in the case of a regenerative satellite is illustrated in figure 7.6.

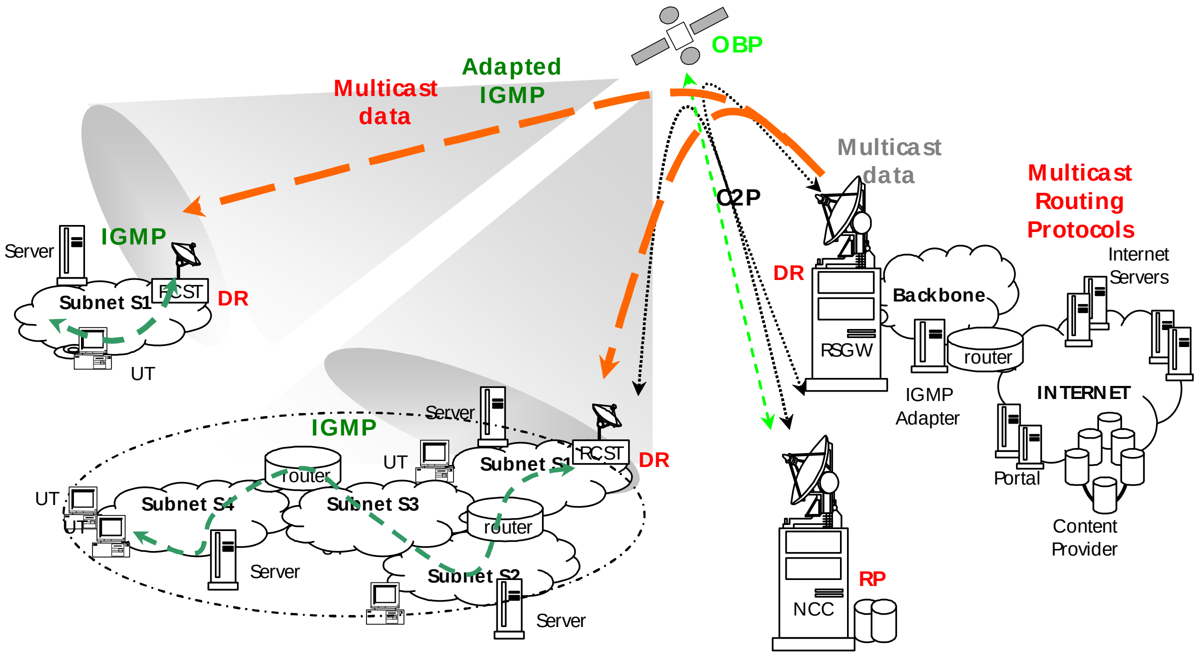

Figure 7.6: C2P & IP multicast interworking

Some additional considerations/suggestions for IP multicast - C2P interworking are provided below:

- Try to keep the IGMP local, between end hosts and their IGMP querier (typically in the RCST), rather than over the satellite, in order to limit the traffic over satellite and minimize the signalling delay.
- Use IGMP over C2P, meaning that the NCC will act as the IGMP querier for the satellite network. C2P messages will help the NCC to manage all the multicast groups present in the satellite network. This implies a higher level of complexity at the NCC, which is the satellite IGMP querier. It allows keeping the controls not only of the multicast resources but also of the multicast groups. In a multibeam regenerative system this will also allow the NCC to build the multicast tree based on the OBP switching tables.
- Use PIM-SM over C2P, allowing each RCST to announce potential multicast sources. This is an extension of C2P functionality, allowing a wider set of functions to establish the multicast tree over satellite. The NCC is informed of any potential multicast source. This is the typical inter-domain multicast scenario.
- Leave C2P only for multicast connection control and perform IGMP adaptation for multicast group management. This solution only allows dynamic multicast group management for star multicast, for multicast sources coming from external networks. For the mesh scenario, this kind of solution will requite a high level of complexity on terminal side, as the terminal should behave as an IGMP proxy or querier on the satellite air interface, depending on the multicast address.

All these different IP multicasting scenarios/solution identified above can be easily identified thanks to a set of multicast flags (table 7.9). These multicast scenarios are based on the reference multicast architecture already defined in [i.9], and each of them identifies a different level of IP dynamic multicasting using a combined solution of S-IGMP and C2P. In all cases C2P is used for connection establishment/release of a multicast connection for IP multicast traffic transmission.

Table 7.9: Example of Multicast Flags

MULTICAST FLAG	DEFINITION	VALUES
Multicast Tx flag	Enables multicast transmission	1 → Tx enabled 0 → Tx disabled
Multicast Rx flag	Enables multicast reception	1 → Rx enabled 0 → Rx disabled
IGMP flag	Enables the IGMP querier functionality of the RCST towards its LAN	1 → Flag enabled 0 → Flag disabled
IGMP proxy flag	Enables dynamic multicast group management based on IGMP adaptation to the satellite environment. This flag can only be enabled if IGMP flag is enabled	1 → Enabled 0 → Disabled
C2P dynamic multicast proxy	Enables dynamic multicast group management by C2P messages (i.e. C2P join/leaves/query/reports)	1 → Enabled 0 → Disabled

The use of the multicast flags in table 7.9 for various multicast scenarios is illustrated in table 7.10.

Table 7.10: Illustration of Multicast Scenarios

Flag/Scenario	Rx	Tx	Igmp	Igmp proxy	C2P proxy	Comment
Static multicast						
Case 1	0	0	0	0	0	No multicast transmission nor reception is allowed
Case 2	1	0	0	0	0	**Static multicast reception.** The terminal can receive a multicast flow by manually configuring the reception of the associated PID
Case 3	0	1	0	0	0	**Static multicast transmission but no reception allowed.** Multicast transmission connections requested via C2P signalling.
Case 4	1	1	0	0	0	**Static multicast transmission and reception enabled.** The terminal can receive a multicast flow by manually configuring the reception of the associated PID. Multicast transmission connections requested via C2P signalling.
Dynamic multicast						
Case 5	1	0	1	0	0	**Dynamic multicast in LAN.** The terminal acts as an IGMP querier. It receives the MMT in the satellite interface and checks the mapping between the desired IP multicast addresses and PIDs. The terminal routes the flows with the associated PIDs to the LAN.
Case 6	1	0	1	1	0	**Dynamic multicast reception both in LAN and at satellite interfaces.** The terminal acts as an IGMP proxy to the satellite network (it sends the IGMP join/leave/query_response messages to the multicast route of the VSN) and IGMP querier to its LAN. Multicast transmission to the satellite network is not allowed.
Case 7	1	1	1	1	0	**Dynamic multicast reception and transmission both in LAN and at satellite interfaces.** The terminal acts as an IGMP proxy to the satellite network (it sends the IGMP join/leave/query_response messages to the multicast route of the VSN) for the reception of the multicast sessions, and as IGMP querier to its LAN. For multicast transmission to the satellite network, on the allowed multicast addressing scope, it acts as a querier towards the satellite network.
Case 8	1	0	1	0	1	**Dynamic multicast in the LAN & multicast transmission not allowed.** For multicast reception all the signalling is sent to the NCC (by C2P messages) that acts both as multicast groups querier and as a controller of the terminals that act as multicast listeners.
Case 9	1	1	1	0	1	**Dynamic multicast in the LAN by IGMP and at the satellite interface by C2P.** For multicast reception all the signalling is directed to the NCC (C2P messages), that acts both as multicast groups querier and as a controller of the terminals that act as multicast listeners. The terminal is allowed to transmit multicast but the querier function is performed by the NCC.
Case 10	1	0	1	1	1	**Dynamic multicast in the LAN and in the satellite interface by S-IGMP; C2P used for OBP multicast switching, no transmission allowed.** This is a multicast scenario only applicable to multibeam satellite networks to include a multibeam OBP dynamic switching. The NCC is not the querier, it only controls which terminals are listening to the multicast flow (C2P join/leave messages) for route management. The terminal injecting the flow in the satellite network will act as IGMP querier. For this terminal transmission is disabled.
Case 11	1	1	1	1	1	**Dynamic multicast in the LAN, and at the satellite interface by S-IGMP; C2P used for OBP multicast switching.** This is a multicast scenario only applicable to multibeam satellite networks to include OBP dynamic switching. The NCC is not the querier, it only controls which terminals are listening to the multicast flow (C2P join/leave messages) for route management. The terminal injecting the flow in the satellite network will act as IGMP querier. In this case, the terminal with this flag configuration may be injecting also multicast groups to the satellite network.

The pure C2P option would be to rely only on C2P messages (join request/response, leave request/response). In this way the NCC is allowed to directly control the multicast sessions as all C2P messages are exchanged between the RCST/RSGWs and the NCC.

In terrestrial networks a router or switch takes the role of querier and periodically requests the hosts in the network to advertise the multicast groups interests. A satellite network in mesh scenarios, by comparison, needs to set up a connection between every pair of hosts connecting in order to support multicast sessions. A direct translation of a multicast protocol to a satellite network implies setting up multiple connections just to transmit the multicast connection signalling information. It might even happen that a host is interested in forwarding a certain group but has no bandwidth available to set up the connection to request it. Or, if there were several feeds, the querier would need to exchange messages with all the feeds, thus multiplying the final number of messages.

All management functions are centralized at the NCC, in particular the resource assignment and the decisions over the routes to be used for the multicast traffic, and also the Query functionality. The NCC acts as a multicast IGMP querier controlling and managing the multicast groups and multicast specific addresses active within the satellite network.

The objective of the multicast protocol is to decide if a flow should or should not be sent through a given branch of a network. In terms of satellite networks this translates to whether to set up or not a point-to-multipoint connection between the feed and the hosts. The NCC is thus heavily involved in the process in all cases. Moving the multicast control towards the NCC and to the control plane substantially reduces the amount of connections and associated messages that are needed, at the price of increased complexity.

A new set of C2P messages can be used to map IGMP messages to C2P. With these messages the following procedures define the pure C2P multicast scenario:

- Add/drop IP multicast addresses of "queried" multicast hosts in the multicast group memberships list contained in the RCST the host is connected to, and in the NCC.
- Add/drop parties in a multicast connection, i.e. aiming at adding/dropping IP multicast addresses of "unsolicited" multicast hosts in the multicast group memberships list contained in the RCST the host is connected to, and in the NCC.

The star dynamic multicast service combining IGMP and C2P allows any RCST that has access to a GW, to join any multicast source from terrestrial networks. The GW will forward into the satellite network the terrestrial network multicast sources under RCSTs' request. The multicast forwarding is dynamic: the IGMPv2 protocol is running between the default GW and RCSTs to discover multicast group membership. A GW only forwards a multicast flow on the uplink if at least one RCST has joined this flow, and the OBP replicates this flow to downlink TDMs covered by the GW.

In this scenario the management of the multicast connections is shared between the NCC and the GW; the NCC exchanges C2P messages with the RCSTs receiving the flow to process the join and leave messages from the RCSTs for the different multicast groups. However, the querier is in this case the GW; IGMP messages will be exchanged between the RCSTs and their associated GW. For that reason this scenario is only applicable to star connectivity and it is not valid for a mesh multicast connection, as the RCSTs can direct their IGMP unsolicited reports to only one destination - their associated GW.

The combined solution (IGMP adaptation and C2P) compared with the pure C2P reduces the complexity on the NCC and improves the responsiveness of the overall network. Another advantage is that there is no delay between the multicast messages and the response is faster, i.e. after a leave message the GW can immediately stop the multicast flow. For IGMP adaptation over the satellite network, IGMP messages are treated like any other data and connections are established to carry them (i.e. connections between the RCST and the GW). If the multicast connection control messages are sent as C2P signalling messages, this will relieve the system from this overhead and will reduce the number of messages exchanged.

One additional advantage of pure C2P multicast solution is that the NCC, being aware of multicast connections, could dynamically inform the OBP of the required switching, optimizing the bandwidth use. The solution is basically a translation of IGMP to C2P messages, therefore the IGMP messages and actor names have been kept to show the parallelism between both. Thanks to this IGMP embedded in C2P, this scenario is ideal for mesh dynamic multicast management.

Annex A (informative): EN 301 790 mesh networking adaptations

This annex contains a selection of extracts from EN 301 790 [i.1]. In particular, the tables are copies of selected tables from EN 301 790 [i.1] as identified in the table titles. Similarly, all of the cross references within these tables should be understood as references within EN 301 790 and not as references within the present TR.

This annex summarizes the adaptations of (amendments to) the EN 301 790 standard to accommodate the needs of mesh networking regarding connectivity and QoS support. These adaptations were first included in version 1.5.1 of EN 301 790 standard. The proposed amendments are intended to be used in association with the Connection Control Protocol (C2P) specification [i.11].

The main adaptation is related to the extension of the Channel_ID from the current 4 bits (16 values) to 8 bits (256 values) or even 12 bits (4 096 values). Other adaptations refer to changes to various tables, descriptors and RL/FL signalling, as a result of Channel_ID extension or in response to additional requirements for C2P support.

The extended Channel_ID will be used to identify both connectivity channels and QoS classes. The connectivity channels are defined to allow an RCST to communicate with other RCSTs in the mesh network or with the Hub/NCC. The QoS classes are consistent with the SatLabs QoS model, developed for star transparent networks. In the case of mesh networks the model applies to each connectivity channel.

The proposed amendments apply to transparent and regenerative network architectures.

If there is any disagreement between this annex and the changes in EN 301 790 [i.1], the published EN 301 790 [i.1] takes precedence.

A.1 New clause on Dynamic Connectivity

Dynamic connectivity is defined in a new clause 4.4 in [i.1], as the capability to establish, modify or release link layer connections between Hub/RCST based upon events occurring on traffic/control or management level.

Dynamically controlled connectivity can be implemented as dynamic connectivity within the existing links in the reference model.

Additional types of satellite interactive networks that exceed the reference model can be built by implementing dynamic connectivity in the form of:

- dynamically controlled connectivity via direct mesh links between RCSTs, through satellite on-board conversion from MF-TDMA to one or more TDM carriers;
- dynamically controlled connectivity via direct mesh links between RCSTs equipped with an MF-TDMA receiver, through the MF-TDMA interaction channel over a transparent satellite.

A.2 Amendments to Common Signalling Channel (CSC) burst format

Common Signalling channel (CSC) bursts are only used by an RCST to identify itself during logon. "Dynamic connectivity" CSC burst data field parameter would allow to signal the support of:

- ETSI C2P TS Protocol Specification [i.11];
- extended Channel_ID (8 bits);
- two Capacity Requests formats (2-byte CR and 3-byte CR);
- C2P signalling by DULM format;
- transparent mesh support.

The Dynamic Connectivity flag is already part of {RCST capability "B"} table in the standard, as shown below.

Table A.1 (EN 301 790 table 1): CSC burst data field parameters

Field Name	Size (bits)	Description/Content
Preamble	variable	Preamble for burst detection and start of burst detection. Definition by TCT. See clause 8.5.5.4. in EN 301 790 [i.1]
RCST Capability "A"	24	See table 2 in EN 301 790 [i.1]
RCST MAC Address	48	RCST MAC address
CSC_Route_ID	16	Enables to define a destination forward (downlink) link for the CSC burst in a regenerative system. If the RCST indicates that this field is overloaded, the value is system dependent
RCST Capability "B"	5	See table A.2
RCST Protocol Version	2	See table A.3
RCST Capability "C"	8	See table A.4
Reserved	8	Reserved
Burst type Identifier	1	"1" (for identification of CSC burst)

Table A.2 (EN 301 790 table 3): RCST capability "B"

Parameter	Bit Size	Description
Dynamic connectivity	1 (b4)	"0" for RCST supporting Dynamic Connectivity according to [i.11] and supporting 8-bit Channel_ID, "1" otherwise. Declaring support for dynamic connectivity may trigger additional, optional logon stages to determine C2P version and capabilities, as defined in [i.11]
Frequency Hopping	1 (b3)	"1" for RCST supporting frequency hopping between adjacent time slots, "0" for RCST requiring one TRF slot between transmissions on different carrier frequencies
DVB-S capability	1 (b2)	RCST capable of using DVB-S on forward link. The field is "1" if the DVB-S is capable, "0" otherwise
DVB-S2 capability	2 (b1-b0)	RCST capable of using DVB-S2 for forward link reception: The field is "11" for not DVB-S2 capable. "01" for DVB-S2 capable of using CCM only, "00" for DVB-S2 capable of both ACM and CCM. Value "10" is reserved

Table A.3 defines the coding used for indication of which version of the DVB-RCS standard that is implemented.

Table A.3 (EN 301 790 table 4): Implemented RCS standard version

Value	Version number
11	Version 1.4 or earlier
10	Version 1.5
01	Reserved
00	Reserved

Table A.4 defines the different bit patterns within RCST capability field "C". The 8-bit field is numbered from LSB to MSB using the notation b0 through b7.

Table A.4 (EN 301 790 table 5): RCST Capability "C"

Parameter	Bit Size	Description
Route_ID_overload	1 (b7)	Indicates alternative, system-dependent use of the CSC_route_ID field. "1" indicates that the field is used as defined in table A.1. "0" indicates that the value is system dependent
Mobility Support	3 (b6-b4)	Mobility support
Continuous ACM	1 (b3)	RCST capable of return link ACM in continuous mode. "0" indicates capable, "1" indicates not capable
NLOS countermeasure support	1 (b2)	RCST capable of non-line-of-sight mobile channel countermeasures. "0" indicates capable, "1" indicates not capable
Transparent mesh reception support	2 (b1-b0)	Indicates support for reception of transparent mesh signals. The field is coded as: 11: No burst mode reception supported 10: Single-carrier burst mode receiver 01: Multi-carrier burst mode receiver 00: Reserved

A.3 Amendments SAC field composition

In [i.1] clause 6.6.1.1, it is proposed to extend the Capacity Request (CR) format from the current 2 bytes to 3 bytes, to allow support of channel identification values higher than 15. An RCST that indicates support of dynamic connectivity (Dynamic Connectivity Flag in CSC set to 0) may support channel identification values higher than 15, and may also specifically indicate such support within the dynamic connectivity control protocol itself.

In the current SAC field (in SYNC bursts and DULM), 3 bits are allocated to the (capacity) TYPE subfield. It is suggested to reduce this subfield to 2 bits and use the remaining bit as CRFORMAT Flag; the value "0" will indicate the current 2-byte CR (CR2), while the value "1" will be for the new 3-byte CR (CR3).

Table A.5: CR2 - original 2-byte format for CR in the SAC field

Width	Value		Min	Max	Unit
1 bit	0-1 "0" "1"	SCALING x1 x16			
1 bit	**0**	**CRFORMAT** 2-byte CR			
2 bits	0-3 "3" "2" "1" "0"	TYPE reserved AVBDC RBDC VBDC			
4 bits	0-15	CHANNEL	0	15	identifier
8 bits	0-255	VALUE	0	255	RBDC: 2-kbps; 32-kbps VBDC/AVBDC: 188 byte or 53 byte; 16 × 188 byte or 16 × 53 byte

Table A.6 shows the format proposed for an extended 3-byte CR (CR3) that could be used as required in the same SAC field. The field is coded to avoid ambiguities when parsing the CRs of the SAC. The extended CR (CR3) has to be supported in mesh networks and by mesh terminals, but is only used when a channel value larger than 15 is needed. When the channel value is 15 or less, CR2 is used, to get space for more CRs in the same SAC.

Table A.6: CR3 – alternative 3-byte format for CR in the SAC field

Width	Value	Description	Min	Max	Unit
1 bit	0-1 "0" "1"	SCALING x1 x16			
1 bit	**1**	**CRFORMAT** 3-byte CR			
2 bits	0-3 "3" "2" "1" "0"	TYPE reserved AVBDC RBDC VBDC			
4 bits		reserved			
8 bits	0-255	CHANNEL	0	255	identifier
8 bits	0-255	VALUE	0	255	RBDC: 2-kbps; 32-kbps VBDC/AVBDC: 188 byte or 53 byte; 16 × 188 byte or 16 × 53 byte

The format of the CR (CR2 or CR3) is first known when decoding the second bit of the first byte (CRFORMAT Flag). Then the following 14 or 22 bits can be decoded properly.

TCT identifies the length of the SAC field available for CR2/CR3, allowing exact fitting of a number of CR2 subfields. Any mix of CR2 and CR3 can be used, as needed. If the SAC field is not filled completely, it is padded according to the following rules:

- Remaining 2 × n bytes are padded by dummy CR2 (VBDC of value=0 for Channel_ID=0, scaling=0).
- Remaining 1 byte in the CR field is skipped by the receiver (will be padded by sender).
- Remaining 2 × n + 1 bytes are split into two sections of 2 × n bytes and 1 byte, taken in sequence, and treated separately according to the above stated rules.

It is assumed that an ATM slot with a 4-byte prefix can take either one or two CR2s or one CR3. When necessary, the SAC field will be padded at the end according to the above rules. ATM with 2-byte prefix cannot support CR3, but only CR2.

The syntax of the SAC field will remain as indicated in table A.7.

Table A.7 (EN 301 790 table 18): syntax of SAC field

Syntax (see note)	No. of bits		Mnemonic information
	Reserved	Information	
SAC_field() {			
if (Route_ID_flag == 0)			
Route_ID		16	uimsbf
if (request_flag == 1)			
for (i=0; i<=capacity_requests_number; i++) {			
Capacity_Request {			
Scaling_Factor		1	bslbf
Capacity_Request_Format		1	bslbf
Capacity_Request_Type		2	bslbf
If (Capacity_Request_Format == 0) {			
Channel_ID		4	uimsbf
Capacity_Request_Value		8	uimsbf
}			
If (Capacity_Request_Format == 1) {			
Channel_ID	4	8	uimsbf
Capacity_Request_Value		8	uimsbf
}			
}			
}			
CR_Pad_Byte		See note	uimsbf
if (M_and_C_flag == 1)			
M_and_C_Message		16	bslbf
if (Group_ID_flag == 1)			
Group_ID		8	uimsbf
if (Logon_ID_flag == 1)			
Logon_ID		16	uimsbf
if (ACM_flag == 0) {			
ACM {			
CNI		8	uimsbf
MODCOD_RQ		8	bslbf
}			
}			
if (MOB_flag ==0)			
Mobility_Control_Message		32	bslbf
Pad_Bytes		see note	uimsbf
}			
NOTE: For SYNC and ATM TRF bursts, the sub-fields used in test statements (Route_ID_Flag, request_flag, capacity_requests_number, M_and_C_flag, Group_ID_flag, Login_ID_flag and ACM_flag) refer to the subfields of the specific timeslot of the TCT for which the RCST has to transmit a SAC field as defined in clause 8.5.5.4. in EN 301 790 [i.1]. For MPEG TRF bursts carrying a SAC field, sub-fields used in test statements refer to the subfields of the SAC_composition of the Adaptation Field Private Data as defined in clause 6.6.1.5. in EN 301 790 [i.1].			

Semantics for the SAC field:

- **Route_ID:** this 16-bit field defines a destination forward (downlink) link that may be used for the prefixed payload in a regenerative system or to indicate a connectivity channel used in association with dynamic connectivity/QoS optimization. Values are system dependent.

- **Capacity_Request:** each capacity request is composed of the following sub-fields:

  - **Scaling_Factor:** this 1-bit sub-field defines the scaling factor of the Capacity_Request_Value sub-field (see table 19).

Table A.8 (EN 301 790 table 19): Scaling_Factor

Value	Scaling factor
0	1
1	16

- **Capacity_Request_Format:** this flag is set to 0 if the 4-bit Channel_ID format is used and to 1 if the 8-bit Channel_ID format is used;
- **Capacity_Request_Type:** this is a 2-bit sub-field specifying the category of capacity request (see table A.9). The capacity categories are described in [i.1] clause 6.8.

Table A.9 (EN 301 790 table 20): Capacity_Request_Type

Capacity_Request_Type Value	Capacity category	Units for Capacity_Request_Value
00	VBDC	Unit of payload size × scaling factor
01	RBDC	Unit of 2kbits/s × scaling factor
10	AVBDC	Unit of payload size × scaling factor
11	Reserved	
NOTE: The payload size is either 53 bytes or 188 bytes according to the encapsulation mode defined at logon.		

- **Channel_ID:** this is a 4-bit field when Capacity_Request_Format=0 and an 8-bit field when Capacity_Request_Format=1. It indicates the channel for which the associated capacity request is being issued. The value 0 000 is the default and indicates that the request is applied to any channel, if not specified otherwise for the system. Other values are system dependent.
- **Capacity_Request_Value:** this 8-bit unsigned integer defines the volume units of payload size or the bit rate in 2 kbits/s of the capacity request as defined in table A.9. A scaling factor as defined in table A.8may be applied.
- **CR_Pad_Byte:** if the space for capacity requests as given by the SAC field specification is larger than what can be utilized, this field contains the necessary number of 8-bit fields with the value "0" to fill up the unused space.

- If the RCST does not have any capacity request to send, it should send a VBDC request with an amount of 0.

A.4 Information Elements (IEs) Range

The Information Elements are as defined in EN 301 790 [i.1], where it has been identified a reserved range of IEs for connection control.

A minimum range of 10 additional IEs will be considered as baseline. A long term solution will require a bigger range.

Table A.10 (EN 301 790 table 26): ATM profile Information Elements

IE type		IE length (see note 1)	IE body
0x00	Capacity Request	2 bytes	As per clause 6.6.1.1 in [i.1]
0x01	M&C	2 bytes	As per clause 6.6.1.1 in [i.1]
0x02	Group_and_Logon_ID	3 bytes	As per clause 6.6.1.1 in [i.1], see note 2
0x03	Message Header	4 bytes	See description below
0x04	Cause	2 bytes	See description below
0x05	Channel_ ID	1 byte	See description below
0x06	Source Address	6 bytes	See description below
0x07	Destination Address	6 bytes	See description below
0x08	Forward Stream Identifier	3 bytes	See description below
0x09	Return Stream Identifier	3 bytes	See description below
0x0A	Type	1 byte	See description below
0x0B	Forward Profile	3 bytes	See description below
0x0C	Return Profile	3 bytes	See description below
0x0D	Security Sign-on Response	8 bytes	As per clause 9.4.9.2 in [i.1]
0x0E	Route_ID	2 bytes	See description below
0x0F – 0x1E	Reserved for connection control	As defined in [i.11]	
0x1F	Wait	As per clause 9.4.9.9	
0x20 – 0x30	Reserved		
0x31	Main Key Exchange Response	As per clause 9.4.9.4	
0x32	Reserved		
0x33	Quick Key Exchange Response	As per clause 9.4.9.6	
0x34	Reserved		
0x35	Explicit Key Exchange Response	As per clause 9.4.9.8	
0x36	ACM	2 bytes	as per clause 6.6.1.1 in [i.1]
0x37	Mobility_Control_Message	4 bytes	as per clause 6.6.1.1 in [i.1]
0x38	Continuous Carrier Control	As per clause 10.4.4 in [i.1]	
0x39 – 0x5F	Reserved		
0x60 – 0x7F	User defined	See note 3	
NOTE 1: "IE length": length (in bytes) of the IE body. NOTE 2: Group_and_Logon_ID: concatenation of the 1-byte Group_ID and the 2-byte Logon_ID fields, in this order. NOTE 3: User defined values can be used for system-specific IEs.			

IE type description:

- Message Header: it identifies the type of message, sets the total length in byte of the message and identifies the connection affected by the connection control signalling.
- Cause: it conveys the reason for the reject of a previous request.
- Channel_ID: defined and utilized according to the present document.
- Source Address: it is the address of the calling end point.
- Destination Address: it is the address of the called end point(s).
- Forward stream identifier: it identifies a single forward information flow pertaining to the connection or to an aggregation of connections; it is either one {VPI, VCI} pair or a single PID, depending on the ATM or MPEG-2 nature of the information flow.
- Return stream identifier: it identifies a single return information flow pertaining to the connection or to an aggregation of connections; it is either one {VPI, VCI} pair or a single PID, depending on the ATM or MPEG-2 nature of the information flow.
- Type: it describes the connection configuration in terms of direction and casting.
- Forward Profile: it describes the priority and the overall amount of resources of the forward streams of the connection.

- Return Profile: it describes the priority and the overall amount of resources of the return streams of the connection.
- Route_ID: see clause 6.6.1.1. in EN 301 790 [i.1].

For DULM with MPEG formatting:

Table A.11 (EN 301 790 table 27): MPEG profile Information Elements

IE type (MPEG)		IE length	IE body
0x00	Capacity Request	2 bytes	As per clause 6.6.1.1 in [i.1]
0x01	M&C	2 bytes	As per clause 6.6.1.1 in [i.1]
0x02	Reserved		
0x03	Message Header	4 bytes	See description below
0x04	Cause	2 bytes	See description below
0x05	Channel_ID	1 byte	See description below
0x06	Source Address	6 bytes	See description below
0x07	Destination Address	6 bytes	See description below
0x08	Forward Stream Identifier	3 bytes	See description below
0x09	Return Stream Identifier	3 bytes	See description below
0x0A	Type	1 byte	See description below
0x0B	Forward Profile	3 bytes	See description below
0x0C	Return Profile	3 bytes	See description below
0x0D	Security Sign-on Response	8 bytes	As per clause 9.4.9.2 in [i.1]
0x0E	Route_ID	2 bytes	See description below
0x0F - 0x10	Reserved		
0x11	Main Key Exchange Response	As per clause 9.4.9.4 in [i.1]	
0x12	Reserved		
0x13	Quick Key Exchange Response	As per clause 9.4.9.6 in [i.1]	
0x14	Reserved		
0x15	Explicit Key Exchange Response	As per clause 9.4.9.8 in [i.1]	
0x16	ACM	2 bytes	as per clause 6.6.1.1 in [i.1]
0x17	Mobility_Control_Message	4 bytes	as per clause 6.6.1.1 in [i.1]
0x18	Continuous Carrier Control	As per clause 10.4.4 in [i.1]	
0x19 - 0x1B	Reserved		
0x1C - 0x1D	User defined	System dependent	
0x1E	Extended IE type for connection control	See description below	
0x1F	Wait	As per clause 9.4.9.9 in [i.1]	

IE type description:

- Message Header: it identifies the type of message, sets the total length in byte of the message and identifies the connection affected by the connection control signalling.
- Cause: it conveys the reason for the reject of a previous request.
- Channel_ID: defined and utilized according to the present document.
- Source Address: it is the address of the calling end point.
- Destination Address: it is the address of the called end point(s).
- Forward stream identifier: it identifies a single forward information flow pertaining to the connection or to an aggregation of connections; it is either one {VPI, VCI} pair or a single PID, depending on the ATM or MPEG-2 nature of the information flow.
- Return stream identifier: it identifies a single return information flow pertaining to the connection or to an aggregation of connections; it is either one {VPI, VCI} pair or a single PID, depending on the ATM or MPEG-2 nature of the information flow.
- Type: it describes the connection configuration in terms of direction and casting.
- Forward Profile: it describes the priority and the overall amount of resources of the forward streams of the connection.

- Return Profile: it describes the priority and the overall amount of resources of the return streams of the connection.
- Route_ID: See clause 6.6.1.1 in EN 301 790 [i.1].
- Extended IE type for connection control: As defined in [i.11].

A.5 Mesh Logon Initialize Descriptor

It is proposed to introduce an additional descriptor that can be used to provide mesh specific information to the mesh capable VSAT at logon. The support of this descriptor is optional. The descriptor format is similar to the format of the Logon Initialize Descriptor, but with the option of including additional information for support of mesh communications. When used, it will be provided instead of the existing Logon Initialize Descriptor as a part of the logon response. The descriptor tag value is 0xB3.

Therefore, after logon, the unicast TIM contains at least the following descriptors: Correction_message_descriptor, either the Logon_initialize_descriptor or the Mesh_Logon_initialize_descriptor, SYNC_assign_descriptor and Satellite_return_link_descriptor.

The descriptor, called Mesh Logon Initialize, can provide the following additional information to the RCST, compared to the Logon_Initialize descriptor:

- the gross maximum CRA assignment that will be allowed for mesh connections;
- the gross maximum number of timeslots per superframe that can be assigned upon AVBDC and VBDC requests for mesh traffic;
- the gross maximum rate that can be assigned upon RBDC requests for mesh traffic resources.

The format of the Mesh Logon Initialize descriptor is provided in table A.12.

This descriptor provides parameters needed for initial logon of a mesh-capable RCST (see table A.12), replacing the logon_initialize_descriptor defined in [i.1] clause 8.5.5.10.4 or such RCSTs. This descriptor is a superset of the Logon Initialize Descriptor, hence it allows the terminal to operate in both star and mesh modes simultaneously.

Table A.12 (EN 301 790 table 77): Mesh Logon Initialize descriptor

Syntax	No. of bits		Mnemonic
	Reserved (see note)	Information	
Mesh_logon_Initialize_descriptor() {			
descriptor_tag		8	uimsbf
descriptor_length		8	uimsbf
Group_ID		8	uimsbf
Logon_ID		16	uimsbf
Continuous_carrier	2	1	bslbf
Security_handshake_required		1	bslbf
Prefix _flag		1	bslbf
Data_unit_labelling_flag		1	bslbf
Mini_slot_flag		1	bslbf
Contention_based_mini_slot_flag		1	bslbf
Mesh_connectivity_flag		1	
Capacity_type_flag		1	bslbf
Traffic_burst_type		1	bslbf
If (Traffic_burst_type == 0) {			
Connectivity		1	bslbf
If (Connectivity == 0) {			
Return_VPI	4	8	uimsbf
Return_VCI		16	uimsbf
}			
Else {			
Return_signalling_VPI	4	8	uimsbf
Return_ signalling_VCI		16	uimsbf
Forward_ signalling_VPI	8	8	uimsbf
Forward_ signalling_VCI		16	uimsbf
}			
}			
Else {			
Return_TRF_PID		13	uimsbf
Return_CTRL_MNGM_PID	3	13	uimsbf
}			
if (Capacity_type_flag == 0) {			
CRA_level		24	uimsbf
VBDC_max	5	11	uimsbf
RBDC_max		24	uimsbf
RBDC_timeout		16	uimsbf
}			
If (Mesh_connectivity_flag == 0) {			
CRA_mesh_level		24	uimsbf
VBDC_mesh_max	5	11	uimsbf
RBDC_mesh_max		24	uimsbf
RBDC_mesh_timeout		16	uimsbf
}			
}			
NOTE: Reserved bits are of type bslbf, and precede the Information bits on the same line. They will be ignored by the RCST. For an encrypted uni-cast TIM, the bit values will be varied in a random manner to avoid encryption spoofing.			

Semantics for the Mesh_Logon_Initialize_descriptor:

- descriptor_tag: the descriptor tag is an 8 bit field which identifies each descriptor;
- descriptor_length: the descriptor length is an 8 bit field specifying the number of bytes of the descriptor immediately following the descriptor_length field;
- Group_ID: this 8 bit field defines which Group ID the terminal is assigned to. This matches the Group_ID used in the TBTP and CMT sections;
- Logon_ID: this 16 bit field identifies the assigned terminal logon identifier, which is used in the TBTP and CMT sections;

- Continuous_carrier: this flag indicates whether the descriptor applies to the MF-TDMA air interface (value "1") or the continuous carrier air interface.(Value "0");
- Security_handshake_required: the value "1" indicates that the security handshake as described in clause 9.4 is to be used. Otherwise it is set to "0";
- Prefix_flag: the value "1" indicates that the Prefix Method according to clause 6.6.1.2 is implemented in the network and the value "0" indicates that it is not implemented;
- Data_unit_labelling_flag: the value "1" indicates that the Data Unit Labelling Method according to clause 6.6.2 is implemented in the network and the value "0" indicates that it is not implemented;
- Mini_slot_flag: the value "1" indicates that the Mini-slot Method according to clause 6.6.1.3 is implemented in the network and the value "0" indicates that it is not implemented;
- Contention_based_mini_slot_flag: the value "1" indicates that the Contention-based Mini-slot Method according to clause 6.6.1.4 is implemented in the network and the value "0" indicates that it is not implemented;
- Mesh_connectivity_flag: a value of "0" indicates that fields defining settings and limits for mesh capacity requests are present. A value of "1" indicates that these fields are absent;
- Capacity_type_flag: a value of "0" indicates that fields defining settings and limits for STAR capacity requests are present. A value of "1" indicates that these fields are absent;
- Traffic_burst_type: this 1 bit field defines the traffic burst type to be used on the return link. The value "0" indicates ATM TRF according to [i.1] clause 6.2.1.1 and the value "1" indicates optional MPEG2-TS TRF according to [i.1] clause 6.2.1.2;
- Connectivity: this 1 bit field defines the connectivity to be used. The value "0" indicates IP connectivity according to [i.1] clause 8.1.1 (Type A RCST) and the value "1" indicates optional ATM connectivity according to [i.1] clause 8.1.2 (Type B RCST). In the case of MPEG2-TS TRF the connectivity is always IP and therefore not signalled;
- Return_VPI, Return_VCI: these fields define the VPI/VCI that the RCST uses in ATM cells on the return link;
- Return_signalling_VPI, Return_signalling_VCI: these fields define the VPI/VCI that is used on the return link for ITU-T Recommendation Q.2931 [i.38] signalling instead of the normal value 0/5. The signalling is used to set up connections for traffic. These parameters can be the same for all RCSTs;
- Forward_signalling_VPI, Forward_signalling_VCI: these fields define the VPI/VCI that is used on the forward link for ITU-T Recommendation Q.2931 [i.38] signalling instead of the normal values 0/5. The signalling is used to set up connections for traffic;
- Return_TRF_PID: this 13 bit field defines the PID that the RCST uses in optional MPEG2 TS packets on the return link for traffic information. This parameter can be the same for all RCSTs;
- Return_CTRL_MNGM_PID: this 13 bit field defines the PID that the RCST uses in optional MPEG2 TS packets on the return link for CTRL/MNGM information. This parameter can be the same for all RCSTs;
- CRA_level: the CRA assignment to the terminal for all STAR communications, in bits/s;
- VBDC_max: the maximum number of payload units per superframe allocated to the RCST through A/VBDC for STAR communications. This is not necessarily an absolute value for each superframe; it can be an average over several superframes;
- RBDC_max: the maximum allowed RBDC data rate, in bits/s for all STAR communications;
- RBDC_timeout: the RBDC timeout, in superframes for all STAR communications. A value of "0" indicates that the timeout is disabled;
- CRA_mesh_level: the CRA assignment to the terminal for all mesh communications, in bits/s;

- VBDC_mesh_max: the maximum number of payload units per superframe allocated to the RCST through A/VBDC for mesh communications. This is not necessarily an absolute value for each superframe; it can be an average over several superframes;
- RBDC_mesh_max: the maximum allowed RBDC data rate, in bits/s for all mesh communications;
- RBDC_mesh_timeout: the RBDC timeout, in superframes for all mesh communications. A value of "0" indicates that the timeout is disabled.

A.6 Network Layer Information Descriptor (NLID) usage

In DVB-RCS standard the Network Layer Information Descriptor is optional and its usage is system-specific. In general it carries SNMP messages in its message body.

In the context of dynamic connectivity the Network layer Information Descriptor can be used to provide the RCST at logon time with information for the establishment of a second default signalling connection, in addition to the default connection that can be established based on information provided in the Logon Initialize Descriptor or Mesh Logon Initialize Descriptor. This is needed for those network scenarios in which two separate logical connections are required for control (e.g. C2P messages) and management (e.g. SNMP messages).

The usage of the NLID for the above purpose requires that the RCST implements SNMP. The Network Layer Info Descriptor is optional, the RCST indicates in the CSC burst if it implements SNMP.

A.7 Timeslot Composition Table (TCT)

The following text is added to the explanation of [i.1] table A.12 Timeslot Composition Table:

- capacity_requests_number: this 3-bit field indicates one less than the number of capacity_requests allowed in the SAC as described in clause 6.6.1.1. In networks using the TRF-prefix signalling method (capacity requests appended to TRF bursts), this field would normally have a value of 0 ("000"). The number of CRs does not need to be indicated for the prefix method. The prefix will be of 4 bytes to allow one 3-byte CR. When using 4 byte prefix only for CR, either one 3-byte CR or two 2-byte CRs will fit in the 4-byte space. If the maximum number of CRs is indicated (different from zero), it defines the maximum number of 2-byte CRs that can be fitted. This field is reserved if the request flag is equal to "0".

A.8 TBTP

The mesh enabled NCC treats the 4 reserved bits + 4-bit Channel_ID field as one 8 bit field and fills the field by (Channel_ID+240) modulo 256.

- NCC will avoid indication of a channel identification value not supported by the RCST (the RCST will not be capable of requesting resources with unsupported channel identification values).
- An RCST compliant with EN 301 790 [i.39] reads only the low nibble as Channel_ID; the high nibble will be expected to be "all ones" and will be discarded.
- The mesh terminal decodes the 8-bit Channel_ID field by adding 16 and doing modulo 256 evaluation.
- NCC should as far as possible avoid assigning channel values above 15 as this will force the RCST to use the less efficient CR3 format instead of the original CR2 format.

Table 28 Terminal Burst Time Plan section is modified as indicated below.

Table A.13 (EN 301 790 table 44): Terminal Burst Time Plane new format

Syntax (see note 2)	No. of bits		Mnemonic
	Reserved (see note 1)	Information	
Terminal_burst_time_plan() {			
SI_private_section_header		64	-
Group_ID		8	uimsbf
Superframe_count		16	uimsbf
frame_loop_count	3	5	uimsbf
for (i=0;i<=frame_loop_count;i++) {			
frame_number	3	5	uimsbf
BTP_loop_count	5	11	uimsbf
for (j=0;j<= BTP_loop_count;j++) {			
Logon_ID		16	uimsbf
Multiple_channels_flag		1	bslbf
Assignment_type		2	bslbf
VBDC_queue_empty_flag		1	bslbf
Start_slot	1	11	uimsbf
If (Multiple_channels_flag == 1)			
If (8-bit Channel_ID is *not* supported)			
Channel_ID	4	4	uimsbf
If (8-bit Channel_ID *is* supported)			
Channel_ID		8	See note 2
Assignment_count}		8	uimsbf
}			
CRC_32		32	rpchof
}			
NOTE 1: Reserved bits are of type bslbf, and will precede the Information bits on the same line. NOTE 2: The use of the 8-bit format of the Channel_ID is restricted to situations where this format is known to be supported by the specific RCST. If this is not explicitly known, the 4-bit format of the Channel_ID will be used towards the specific RCST.			

- Channel_ID: this 4-bit or 8-bit field indicates the channel to which timeslots are being assigned. Values are managed by the NCC. This field is present if the multiple_channels_flag is set to "1". The 4-bit field is used if the 8-bit field is not supported; it carries Channel_ID values in the range 0-15. The 8-bit field is used if it is supported; it is a uimsbf field that carries Channel_ID values in the range 0-255, coded as {(Channel_ID + 240) modulo 256}. Support for the 8-bit Channel_ID is indicated at logon time in the CSC burst, as defined in [i.1] clause 6.2.3.

A.9 Implementation type descriptor

This descriptor , as described in table A.14, supports indication of the generic implementation type of the hub. The indicated options can be interpreted by the RCST without additional information. The information provided can assist the RCST in choosing a way to operate that will work. Relevant to C2P, it will include the C2P protocol version. The descriptor supports system specific extensions. The descriptor tag value is 0xB5.

Table A.14 (EN 301 790 table 78): implementation type descriptor

Syntax	No. of bits		Information
	Reserved (see note)	Information	
Implementation_type_descriptor() {			
descriptor_tag		8	uimsbf
descriptor_length		8	uimsbf
ncc_protocol version		8	uimsbf
location_update_allowed_flag	2	1	bslbf
rbdc_accepted_flag		1	bslbf
vbdc_accepted_flag		1	bslbf
avdbc_accepted_flag		1	bslbf
timing_offset_flag		1	bslbf
timing_reference_flag		1	bslbf
C2P_protocol_version	5	3	uimsbf
hub_type_id		24	uimsbf
hub_sw_id		24	uimsbf
user_options_count		8	uimsbf
For(i=0; i < user_options_count; i++) {			
user_options_byte		8	uimsbf
}			
For(i=0; i < n; i++) {			
reserved_byte		8	uimsbf
}			
}			
NOTE: Reserved bits are of type bslbf, and precede the Information bits on the same line. They will be ignored by the RCST. For an encrypted uni-cast TIM, the bit values will be varied in a random manner to avoid encryption spoofing.			

Semantics for the Implementation_type_descriptor:

- descriptor_tag: the descriptor tag is an 8 bit field which identifies each descriptor;
- descriptor_length: the descriptor length is an 8 bit field specifying the number of bytes of the descriptor immediately following the descriptor_length field. This descriptor is specifically foreseen to be extended at the end with more specific bytes in future revisions, to provide space for more information about the generic implementation type. If the descriptor_length specifies fewer bytes than is needed to hold the information specified in table 78 from EN301 790 [i.1], the descriptor is truncated at the corresponding point. This way, the information can be safely interpreted from the first byte following the header with the interpretation known by the specific RCST, skipping any unknown content. A specific system may use a short or long descriptor to convey as much information as desired. However, if a non-zero value of user_options_count is specified, the descriptor_length will be sufficient to accommodate the specified number of user_options_bytes;
- ncc_protocol_version: this 8 bit field indicates the RCS protocol version implemented by the NCC, coded as specified in table 26.

Table A.15 (EN 301 790 table 79): standard version implemented by the NCC

Value	Version number
2-255	Reserved
1	Version 1.5
0	Version 1.4 or earlier version

- location_update_allowed_flag: a "1" indicates that the hub generally allows location updates issued by the RCST in the CSC slot as specified in clause "CSC burst format". A "0" indicates that location update by use of the CSC slot is generally prohibited, but may be specifically allowed for the RCST in question according to system specific methods;
- rbdc_accepted_flag: a "1" indicates that the hub honours RBDC requests on the default channel. A "0" indicates that the hub may unconditionally discard RBDC requests on the default channel;

- vbdc_accepted_flag: a "1" indicates that the hub honours VBDC requests on the default channel. A "0" indicates that the hub may unconditionally discard VBDC requests on the default channel;
- avbdc_accepted_flag: a "1" indicates that the hub honours AVBDC requests on the default channel. A "0" indicates that the hub may unconditionally discard AVBDC requests on the default channel;
- timing_offset_flag: a "1" indicates that the hub requires that the RCST offsets its timing from the native NCR packet source reference point as indicated in the NCR payload field. A "0" indicates that an RCST may attempt to log on using the native timing of the system even if it discards the optional NCR payload;
- timing_reference_flag: a "1" indicates that the hub applies the nominal position of the return link satellite for each return path as the native NCR packet source reference point, referring to the point of time when the PCR packet with the NCR value entered the channel interleaver and FEC encoder at the nominal link rate, or to the applicable start of the frame signal (SOF) as specified for DVB-S2 ACM/VCM. A "0" indicates that the hub may apply a system specific native NCR packet source reference point;
- C2P_Protocol_version: this 3-bit field defines the version of the connection control protocol defined in [i.11] that is supported by the NCC. The coding of the field is defined in table A.16.

Table A.16 (EN 301 790 table 80): connection control protocol version implemented by the NCC

Value	Version number
0	No connection control protocol supported
1	Version 1
2 to 7	Reserved

- hub_type_id: a 24 bit parameter that is used to identify the type of hub. The value used can be an OUI value registered by the hub vendor in the IEEE Registration Authority as a company id. If such an OUI is not applicable for the hub, the field will be set to 0xFFFFFF to unconditionally indicate an unknown hub type;
- hub_sw_id: a 24 bit parameter that identifies the SW running on the hub in the context of a known hub type. It should be discarded if the hub_type_id is unknown;
- user_options_count: the number of bytes used to indicate user defined options;
- user_option_byte: the concatenated user_option_byte field contains user defined options. These options will be interpreted in the context of a known hub_type_id, and may have to be interpreted in the context of the hub_sw_id. The content will be discarded if the hub_type_id is unknown;
- reserved_byte: this content will be discarded by the receiver.

List of figures

List of tables

History

Document history		
V1.1.1	January 2009	Publication